



Article

# Designing Behavioural Artificial Intelligence to Record, Assess and Evaluate Human Behaviour

Hanno Hildmann

Department de Ingeniería de Sistemas y Automática, Universidad Carlos III de Madrid (UC3M),  
Av. Universidad, 30, 28911 Légnés, Spain; hanno@cypherpunx.org

Received: 13 March 2018; Accepted: 28 August 2018; Published: 25 September 2018



**Abstract:** The context of the work presented in this article is the assessment and automated evaluation of human behaviour. To facilitate this, a formalism is presented which is unambiguous as well as such that it can be implemented and interpreted in an automated manner. In the greater scheme of things, comparable behaviour evaluation requires comparable assessment scenarios and, to this end, computer games are considered as controllable and abstract environments. Within this context, a model for behavioural AI is presented which was designed around the objectives of: (a) being able to play rationally; (b) adhering to formally stated behaviour preferences; and (c) ensuring that very specific circumstances can be forced to arise within a game. The presented work is based on established models from the field of behavioural psychology, formal logic as well as approaches from game theory and related fields. The suggested model for behavioural AI has been used to implement and test a game, as well as AI players that exhibit specific behavioural preferences. The overall aim of this article is to enable the readers to design their own AI implementation, using the formalisms and models they prefer and to a level of complexity they desire.

**Keywords:** behavioural artificial intelligence; propositional logic; modal logic; behavioural psychology; rational choice theory; utility theory; game theory; computer games

## 1. Introduction and Outline

The philosophical question whether a machine can be intelligent is “as old as computers themselves” [1] and Alan Turing’s 1950 article on “Computing Machinery and Intelligence” [2] famously opens with the question “Can machines think?” Some of the most widely known successes of Artificial Intelligence have been achieved by game-playing programs, which, one by one, have dismantled strongholds of human intelligence [3]. Chess, for centuries considered a pinnacle of human intelligence and intellect, was famously conquered by the program Deep Blue when it defeated the reigning human world champion in 1997 [4]. Recently, the best living human player of the game of Go, a game orders of magnitude more complex than Chess, was beaten by AlphaGo [5,6] (in a sweep victory [7]). In 2017, an improved version of that program [8] autonomously learned Go, Chess and other games, and then proceeded to beat the best existing players (which at this point are all machines). Machines have blown away the best human players of Jeopardy! [9] (a natural-language based TV game-show that relies on the understanding of subtle hints, jokes and riddles) as well as outperformed humans in “Heads-up limit hold” Poker [10] (where the analysis of the opponents’ playing behaviour is crucial) [11,12]. With regard to traditional board- and card-games, the age of man has truly come to an end.

Machines have started to outperform humans in virtually all areas of *rational* decision making. Machines can *do the right thing*. However, when it comes to making bad, irrational or obviously false decisions, most of our behavioural models struggle or fail. Achieving realistic and human-like machine behaviour seems to be far more difficult than building programs that outsmart us.

### 1.1. Motivation

The ability to exhibit natural behaviour patterns (however these may be defined) on a level that is compatible with human behaviour is considered a key challenge for robots, virtual agents and intelligent machines designed to interact with humans. Recent years have seen a massive increase in the use of intelligent interaction partners designed to assist humans in virtually all areas of our daily lives. While it may not be a fundamental requirement for all application areas, the ability to elicit social interaction or to appropriately respond to specific human behaviour patterns undoubtedly has the potential to improve performance of a system as well as greatly increase its acceptance by society.

### 1.2. Context

Our approach to designing a behavioural Artificial Intelligence, presented in Section 5.6, is firmly rooted in the desire to design an automated tool for the assessment, evaluation and comparison of human behaviour (to which the entire Section 5 is dedicated). Such an undertaking requires (among other things) a formalism (cf. Section 3 for an overview of TACT (Target, Action, Context and Time) and logic, and Section 5 for how these are combined into a tailored formalism) to express behaviour. In addition, to guarantee comparable results, it is necessary to ensure that conditions surrounding the assessment are comparable and entirely controlled by the researcher. Relatively simple computer games are ideally suited to facilitate the latter and Section 4 provides our formal definitions for such games. The idea of using computer simulations for psychological research is not new [13]. Herbert Simon observed [14] that *“An important part of the history of the social sciences over the past 100 years, and of their prospects for the future, can be written in terms of advances in the tools for empirical observation and in the growing bodies of data produced by those tools”*. The origin of the approach and model presented in this article stems from our work on the design and implementation of tools specifically designed to be used in an automated manner. The larger is the scope of the empirical study for which such a tool will be used, the more important it is to use unambiguous formalisms and to ensure that they can be fully automated. The integrity of the collected data and the representative nature of the study will depend on this and, by extension, so will the usefulness of the tool itself.

### 1.3. Outline

In Section 2.4, we discuss a model for human behaviour used in the field of behavioural science, specifically behavioural psychology (Section 2). Rationality and intelligence both play a part in these fields. In Section 3, we introduce formalisms from psychology, logic, and game theory and, in Section 4, we argue that computer games are commonly used in psychology and education and that certain games can provide a discrete, well-defined and controllable environment to assess and evaluate behaviour. In other words, computer games can be well-suited as the embodiment of the proposed formal approach. The second half of Section 4 provides formal definitions and a model for such games. Using these and the formalism discussed in the section at hand, we discuss in Section 5 which aspects of behaviour we want to express formally, how such formal statement can be evaluated to be true or false (in an efficient and automated manner) and, ultimately, how the ability to evaluate formally stated behavioural preferences against a set of behavioural choices can be used to drive AI behaviour under rational choice.

### 1.4. Disclaimer

This is *not* an article from the field of psychology using computing science technologies; instead, it is a computing science article aiming to make a contribution to the fields of behavioural psychology and AI. The author is by no means an expert in the field of psychology; the model used was chosen on the basis of it being widely used by practitioners. The presented approach is designed with a chosen model as foundation, but it is not at all model-dependant (which is a good thing).

## 2. Behavioural Sciences

Behavioural science is a conglomerate of many disciplines, all of which are dedicated in one way or another to the exploration of actions and interactions between humans, animals or, generally, entities. For example, social cognitive neuroscience attempts to identify and describe the brain areas and mechanisms that mediate social life [15] while behavioural psychology is a branch of psychology that focuses on observable behaviours. Such behaviours range from collaboration to solve a problem or a task in the physical world [1] over rational choice (or the absence thereof) [16] and risk taking [17] to language acquisition and evolution. This section draws on previous publications [18–24].

In this section:

- (a) We elaborate on the claim that experts in psychology disagree on what affects behaviour (Section 2.1).
- (b) We argue that game theoretic models regularly fail to correctly predict human behaviour (Section 2.2).
- (c) We discuss the dilemma of competitive versus cooperative behaviour (Section 2.3) which our proof-of-concept behavioural artificial intelligence (cf. Section 6) was designed to face.
- (d) We introduce (Section 2.4) the model from behavioural psychology used for our formalization (Section 3).

### 2.1. Behavioural Psychology

Most people would readily agree that our decisions regarding action and behaviour are partly controlled by their anticipated outcome (cf. the model in Section 2.4). Independent of someone's values, we humans model and reinforce what we value and our actions reflect these values [25]. However, this is not consistent with human behaviour as observed, e.g., in gambling, drug addiction and health care.

Edward Thorndike and B. F. Skinner considered behavioural learning a matter of reinforcement [26–28]. According to them, organisms respond to reinforcement and, when found in previously experienced situations and presented with known incentives, will react the same way [29]. The theory that the learning of behaviour is achieved through repetition and reinforcement has many critics (e.g., [30]), which are quick to point out that this theory completely removes mental processes and models from the conscious decision. While these theories might find justification when investigating (and observing) animals with lower intelligence, they leave no room for the subjective nature of human behaviour. As stated in [31], Skinner said that people act because of conditions that make them act. Chomsky, who also has his critics ([32]), argued that certain important questions are ignored in this approach [33] and stated that *"[o]ne would naturally expect that prediction of the behavior of a complex organism (or machine) would require, in addition to information about external stimulation, knowledge of the internal structure of the organism, the ways in which it processes input information and organizes its own behavior"* [34]. The underlying organization of behaviour, and within it the preferences for some behaviours over others, are (according to Chomsky) important factors in its analysis. Our model for behavioural artificial intelligence (cf. Section 5.6) also uses a preference function to order behaviours.

A clear contrast to behavioursim is the approach of Carl Rogers [31] who believes that people act on their internal state, i.e., they act because of states they feel (and not, as Skinner postulated, because of the external conditions that make them feel like this). According to Rogers, humans inherently have the ability to *"figure out what is best for them"*, which is reflected by his view that therapy is essentially a learning experience and not, as Freud saw it, a *battleground* of conflicting drives [31].

In the real world, our behaviour is generally bounded by time and our decision making processes have evolved to address this. Daniel Kahneman [35] argued that there are several processes in our brain that operate at different speeds as well as levels of complexity and rationality. Gigerenzer and Goldstein [36] proposed that the brain operates a number of simple heuristics [37] similar to a *toolbox* [38] from which we use simple tools adaptively as we encounter decisions in life. In the field of Artificial Intelligence, Minsky [39] proposed in 1988 that cognition and intelligence are the product of a number of specialized *agents*, collectively forming what he referred to as a *society of mind*.

## 2.2. Game Theory and Rational Choice

Game theory has provided the field of behavioural sciences with a variety of rigorous models aiming to predict and understand decision making. However, as, e.g., Sanfey [40] pointed out, humans live in a highly complex social environment and there are a host of results that show that humans do not behave rationally (as defined in game theory). As a matter of fact, game theoretic models often fail to adequately predict human behaviour. Specifically, participants in experiments often fail to choose a so-called *Nash-equilibrium* [41,42], i.e., a state in which a player can expect to get the highest payoff assuming that the other players also base their decision on maximizing their subjective payoff. As it turns out, human decisions are often not based on selfish considerations and social factors are frequently considered [40]. Milton Friedman wrote [43] that “[e]conomics as a positive science is a body of tentatively accepted generalizations about economic phenomena that can be used to predict the consequences of changes in circumstances” and Shapley [44] portrayed economics as the most successful social science. This claim might find some traction because the generalisations that are made are about human behaviour, and the observed behaviour is likely to be representative because humans are inherently greedy, and (by *traditional* economic theory [16]) can be expected to behave *rationally* when it comes to economic (social) choice.

Sanfey [40] reported on investigations into social choice: “the study of decision-making attempts to understand our fundamental ability to process multiple alternatives and to choose an optimal course of action”. The author continues to say that this ability has been the matter of many studies in a number of fields and that these studies used a variety of theoretical assumptions as well as measurement techniques, and mentions the discipline of game theory as being one of the fields that has greatly contributed to these investigations. Both rational decision and social choice are relevant for what we discuss in this article: the context of the presented work is to provide a tool for the investigation of human behaviour, and, to achieve this, a game-playing behavioural AI driven by the same formalism was designed. Specifically, we use Simon’s model for rational choice, *subjective expected utility (SEU) theory* (explained in Section 3.3.3) is used in Section 5.6 (p. 25) for our model of rational behaviour in AI players. The connection between game theory and rational choice is discussed in Section 3.3. (p. 9).

## 2.3. Cooperative/Competitive Behaviour

As is evident from the large body of literature, cooperative (and thus also competitive) behaviour is amongst the behaviours that have been greatly investigated in behavioural sciences. Some consider it to be a key aspect of human behaviour [45,46]. Simon mentioned [14] altruism versus egoism, and the potential usefulness of cooperation in a Darwinistic system. Humans, as a social species, have achieved a lot through cooperative behaviour and coordinated actions with others [47] and one would expect to find similarities everywhere. However, this cannot be generalized as different societies can vary greatly with regard to cognitive stances [17]. For example, linguistics: in English, speakers take a different cognitive stance (egocentric) than speakers from certain non-Western cultures (allocentric).

Several theories have been proposed to explain the evolution of human cooperation [48]. In the literature, hypothetical test scenarios have been used to investigate test subjects’ decisions with respect to acting cooperatively or competitively. Since it is regularly [49] used to explain a *Nash Equilibrium* and to illustrate the difficulties of making a rational cooperative decision [50], we briefly explain *The Prisoner’s Dilemma* [14,41,51]. In this well-known conundrum, only two choices are offered in a hypothetical situation: cooperation and defection. The setting is that two players are both presented with these choices and are advised that their choices will affect the outcome for themselves as well as for the other player, as shown in Table 1 below:

**Table 1.** The payoffs for the various combinations of choices in The Prisoners Dilemma.

	Cooperate Player 1	Defect Player 1
Cooperate Player 2	(B, B)	(A, D)
Defect Player 2	(D, A)	(C, C)

The entries in the table represent outcomes of the game, (X, Y) being the outcome for player 1 (X) and player 2 (Y). The values A, B, C and D are related to each other as follows:  $A > B > C > D$  and such that  $B > \frac{A+D}{2}$ , where “ $X > Y$ ” means that X is preferable over Y. Everything about this game is known to both players, except the other player’s choice.

This dilemma is a thought experiment that can be traced back to Thomas Hobbes and Jean-Jacques Rousseau [52]. It shows the problem of achieving mutual cooperation [50] in a scenario where it is possible to gain an advantage over the other player. The argument is that the most rational choice is to defect, as this will either yield the best possible result (in the case when the opponent cooperated) or (if the opponent also defected) at least avoids the worst possible outcome. It was studied during the cold war [52] by members of the RAND corporation [53] and was used by von Neumann to advocate a nuclear first strike against Russia [54].

The problem was extended by considering a finite number of these choices between the same two players, called the *iterated Prisoner’s Dilemma* [55], which is often described as a game-theoretic paradigm for the evolution of cooperation based on reciprocity [56]. Humans exhibit (sometimes highly) irrational behaviour by, e.g., adapting a much more forgiving approach when there is the opportunity for reciprocity. Due to this, the prisoner’s dilemma has been used “to show how altruism can develop in animal communities, including human societies” [52]. When faced with a repeated choice as in the iterated prisoner’s dilemma, a common strategy called “*tit for tat*” (“tip for tap” meaning equivalent retaliation i.e., where one player will take the stance taken by the other player in the previous game), emerges. The strategy goes as follows: start in the initial round by cooperating and in every subsequent round adopt the decision taken by your opponent in the previous round. This has been investigated extensively (e.g., [57,58]).

#### 2.4. Modelling Behaviour

Simon’s model [16] of rational choice, subjective expected utility theory (SEU), discussed in Section 3.3, is based on the assumption that people are perfect rational decision makers [59]. Clearly, this is not the case and there is a host of work in the literature making this point [60]. Prospect theory [61] adds to this by showing that humans are not even consistent in their evaluation of specific outcomes of actions. As a result, even if there were a single model for rational decision making, it would be subject to changing interpretations. Daniel Kahneman suggested that our inability to be equally aware of the bad things as we are of positive and good things might be a natural human trait [62].

In the context of computer games and artificial intelligence, *player modelling* has been an important research domain for decades where one of the key challenges is behavioural consistency [63]. There is a growing number of complementary and competing models [64], some of which have been strongly influenced by various aspects of human emotional behaviour [63]. Generally, player modelling can be seen as a loose concept [65] aiming at the design and study of computational models for players in games [66], be it human or AI players. For the industry, predicting player behaviour (with the aim to adjust the game to the player) [67] or the fact that behavioural and cognitive capabilities akin to that of humans has the potential to greatly increase the believability of agents [66] have been a driving factor. The view that certain social and cognitive aspects are basically required by any intelligent system [64] is supported by experts from the field of AI, such as Marvin Minsky [39] and Herbert Simon [68].

As stated in [18], attitude is a *hypothetical construct* that represents an individual mental predisposition either for or against some concept or idea. Simply changing the framing of a question or problem can shift our preferences for specific outcomes [69]. This established fact is incompatible with



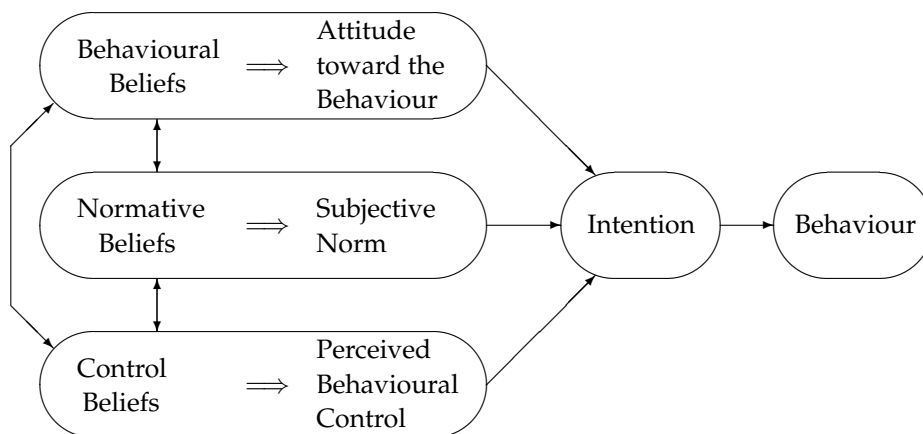
many models for rationality and standard economic accounts and the neurobiological mechanisms causing it are not yet understood [70]. In the context of this article and given that we are not aiming to compete with the different models in the field, we suggest focussing on the attitude of subjects towards some aspect (e.g., environmental issues) instead of the more commonly practised approach of attempting to have the subjects report on their behaviour directly.

The Theory of Planned Behaviour (ToPB) in psychology, a theory regarding the link between attitudes and behaviour ([71,72]), provides a model for human behaviour that treats the attitude a person has towards some action as a relevant factor in the decision making process as to whether to execute this action. This theory implies that the attitude someone has towards a certain behaviour influences that person's likeliness to subsequently exhibit it. It is of interest to us in the context of this article, as we used the related description for behaviour as the basis for our formalism.

According to ToPB and with respect to actions and behaviour, human decision making is guided by three different considerations and beliefs:

1. Behavioural beliefs: Behavioural beliefs are someone's expectations about the likely outcome of actions, paired with the subjective view of these outcomes.
2. Normative beliefs: Normative beliefs are the opinion of others regarding the outcomes of actions, the personal intention to adhere to these peer standards as well as the desire of the individual to live up to the expectations of one's peers.
3. Control beliefs: Control beliefs are one's level of confidence that they have control over all relevant factors required to bring about an outcome.

Figure 1 illustrates this model and how the mentioned beliefs and considerations influence one's intentions and subsequently one's behaviour. This theory has formed the basis for previous work on the evaluation and the assessment of human game-playing behaviour.



**Figure 1.** The Theory of Planned Behaviour (simplified illustration) [71,72].

### 3. Formalisms

In this section, we introduce established models and formalisms from:

- (a) behavioural psychology (i.e., the TACT paradigm, Section 3.1) to formally describe behaviour;
- (b) philosophy and logic (i.e., classical propositional logic as well as modal logic, Section 3.2); and
- (c) game theory (i.e., game theory, utility theory and subjective expected utility theory, Section 3.3).

The first provides a formalism to describe behaviour, the second a formalism to express the so-described behaviour in a way that the evaluation thereof can be automated, and the third models for rational decision making that can be augmented to include behavioural preferences provided this way.

### 3.1. The TACT Paradigm

With respect to observable aspects of behaviour, we continue to use the works of Icek Ajzen as our reference point, specifically the TACT (Target, Action, Context and Time) paradigm that was suggested for the design and the evaluation of questionnaires (within the context of ToPB related research). We previously discussed our choice for TACT in [19–21,24].

Ajzen [73] argued that, to define behaviour sufficiently (within the context of ToPB), four aspects of any behaviour have to be defined: Target, Action, Context and Time (TACT). His running example is “walking on a treadmill in a physical fitness center for at least 30 min each day in the forthcoming month”, and he states that the distinction between these four aspects is not always clear. Ajzen himself pointed out this ambiguity and suggests that there are many possible additions to his basic TACT paradigm (e.g., “within the next month” can include “next Tuesday”). We argue that, within the scope of this article, the four TACT aspects suffice. However, we point out that we neither expect these four nor the TACT approach itself to be the optimal solution to all applications or problems. There are a number of theories in the field of behavioural psychology and the adopted theory and paradigm will depend on the specific focus of the application. The specifics of the project for which the formalism (presented in Section 5) is eventually used will determine the extent to which a finer grained distinction if TACT (or indeed a different paradigm) is required. Furthermore, complicated extensions will complicate the formalism presented in this article without adding value to the conceptual approach and are therefore omitted here. For most applications, the presented level will suffice; for more demanding applications, the extensions would probably have to be very specific to the application, but would not be conceptually different. Therefore, we argue that the extent of the introduced material suffices for this article.

#### Examples

As above, our running example [73] is “walking on a treadmill in a physical fitness center for at least 30 min each day in the forthcoming month”. The TACT elements in this and similar statements could be:

Action: walking, exercising, working out

Target: on a treadmill, on a stair master, on a walking machine

Context: at home, in a physical fitness centre, in the gym

Time: for at least 30 min each day in the forthcoming month

Another set of examples, closer to what we use (see Section 5.2), is:

- (1) This summer, Person A sells a boat to Person B in Edinburgh.
- (2) Last year Person A bought a boat from Person C in Glasgow.
- (3) The next five months Person A will not sail on the Clyde.

The TACT elements in these examples could then be identified as:

Action: buy, sell, sail

Target: Person A, Person B, Person C

Context: in Glasgow, in Edinburgh, on the Clyde

Time: this summer, last year, the next five months

We elaborate on this in detail in Section 5.2 and continue first by introducing logic as a formal and well-defined language which can be evaluated by automated processes.

### 3.2. Propositional and Modal Logic

In the most general description, logic “may be regarded as the systematic study of thought” [74] and it has been referred to as the “theory of good reasoning” [75]. “Thought” and “reasoning” are concepts associated with the field of psychology [76], which also concerns itself with “decisions” and

“choices” [77]. However, we should not treat the field of logic as a sub-category of psychology (which it could, arguably, be considered [78]), because the *logical distinction between valid and invalid inference does not refer to the way we think* [79]. Kant wrote that the emphasis of logic was not on *how* we think but on *how we ought* to think [80], and, indeed, humans have been observed to think and reason in an irrational and very un-logical way [81]. The use of logic in this article is exclusively as a formalism, *to make unambiguous statements about behaviour*. We discuss common models for behaviour from the field of behavioural psychology, and, when it comes *to investigating behaviour*, it is those we draw upon.

### 3.2.1. Propositional Logic

Propositional logic (PL) (cf. [82]) is not concerned with anything *but* propositions which are statements that are either true or false [83]. As such, PL treats the world about which it reasons as a snapshot, a photograph if you will, which is static. Therefore, PL is incapable of expressing change, uncertainty or probabilities but only individual statements  $p$  and  $q$  and their truth values (so named by Gottlob Frege [84]). These propositions could mean anything; a logician is not really concerned with the actual state of affairs as long as there are  $p$ s and  $q$ s to reason about.

This implies a very narrow minded view on the world as it excludes anything that cannot be unambiguously evaluated. However, for our purpose, this view is sufficient. We consider a finite set of such propositions, traditionally called  $\Phi$  [85], to be enough for an adequate description of the world.

#### Syntax and Semantics

In what follows, we denote such atomic statements in  $\Phi$  by  $p$  and  $q$  and introduce the means to combine them into more complex statements. We now introduce two such connectors: the first, “ $\neg$ ” we call *not* (usage  $\neg p$ ) and “ $\wedge$ ” which we call *and* (usage  $(p \wedge q)$ ). Since we take this narrow view on the world where everything is either true (T) or false (F) (the latter being defined as  $\neg T$ ), we can define the meaning of these connectors unambiguously by giving their *truth tables* [86], i.e., by defining their truth values for all possible combinations of truth values of the propositions they cover (cf. Table 2). In the first case, there are only two ( $p$  can only be true or false), while, in the latter, there are four: both  $p$  and  $q$  are true,  $p$  is true and  $q$  is not,  $p$  is false and  $q$  is true, and both  $p$  and  $q$  are false:

**Table 2.** The truth tables defining the  $\neg$  (left) and the  $\wedge$  (right) operator.

$p$	$\neg p$	$p$	$q$	$(p \wedge q)$
T	F	T	T	T
F	T	T	F	F
		F	T	F
		F	F	F

The syntax for any *well formed formula* (wff)  $\phi$  in  $PL(\Phi)$  (propositional logic formulae constructed over propositions contained in  $\Phi$ ) is defined as follows:  $\phi \in PL(\Phi)$  if and only if (iff):

$$\phi = p \mid T \mid \neg\psi_1 \mid \psi_1 \wedge \psi_2$$

with  $p \in \Phi$  and  $\psi_1, \psi_2$  also wff of  $PL(\Phi)$ . The above means that a formula  $\phi$  is: (a) a proposition; (b) a truth value; (c) the negation of a formula; or (d) a conjunction of two formulae.

#### Syntactic Sugar

We use symbols which enhance readability of formulae but do not add expressiveness to the language. These symbols can be considered mere abbreviations for longer formulae which would otherwise be confusing or counter-intuitive; the above introduced “false” (F), defined as  $\neg T$ , is an instance of such *syntactic sugar*. Commonly, PL has additional symbols for three more operators:



$\vee$  stands for “or”,  $\rightarrow$  for “if ... then ...” and finally  $\leftrightarrow$  represents “if, and only if”. Their syntax is as follows:  $(p \vee q)$ ,  $(p \rightarrow q)$  and  $(p \leftrightarrow q)$ . However, a well-known result, *DeMorgan’s Law* [87], states that these connectives are merely abbreviations and that they can be rewritten using only the operators  $\wedge$  and  $\neg$ :

1.  $(p \vee q)$  is equivalent to  $\neg(\neg p \wedge \neg q)$
2.  $(p \rightarrow q)$  is equivalent to  $\neg p \vee q$
3.  $(p \leftrightarrow q)$  is equivalent to  $(p \rightarrow q) \wedge (q \rightarrow p)$

These rewritten rules come in handy in Section 5.4 when we provide a mechanism to evaluate a statement in a specific state  $s_j$  of a game. The remaining ambiguities can be removed by introducing brackets to distinguish, e.g., between  $p \vee (q \wedge r)$  and  $(p \vee q) \wedge r$ . With this in place, translating these formal statements into some semi-natural language  $\mathcal{L}$  and vice versa is straightforward (see Section 5.3).

### 3.2.2. Modal Logic

Generally speaking, modal logic is any logic with *modalities*. Recall that PL is only concerned with what is true or false. This, however, is very far from modelling a variety of issues of our daily lives where we constantly encounter things that are *possible*, *probable*, or sometimes *follow necessarily*. Modal logic allows for models that entertain two different, possibly even contradicting, states that the world could be in. Such states constantly occur in our lives and as our world around us evolves we continuously update and change our representation thereof. In other words, change and uncertainties are part of our environment and a formalism meant to describe behaviour will have to allow for this.

#### Syntax

Modal logic is propositional logic enriched with modalities. Syntactically this is achieved by adding a unary modal operator,  $\Diamond$ . The syntax for any *well formed formula* (wff)  $\phi$  in  $ML(\Phi)$  (modal logic formulae constructed over propositions contained in  $\Phi$ ) is defined as follows:  $\phi \in ML(\Phi)$  iff:

$$\phi = p \mid \top \mid \neg\psi_1 \mid \psi_1 \wedge \psi_2 \mid \Diamond\psi_1$$

with  $p \in \Phi$  and  $\psi_1, \psi_2$  also wff of  $ML(\Phi)$ . The above means that any well formed formula of modal logic is: (a) a proposition; (b) a truth value; (c) the negation of a formula (of propositional or modal logic); (d) a conjunction of two formulae (of propositional or modal logic); or (e) a formula of propositional or modal logic with a modality added. We add a second symbol for the dual,  $\Box\phi$ , which is defined as:  $\Box\phi = \neg\Diamond\neg\phi$ . The reason this is (again) syntactic sugar is explained below.

#### Semantics

In *Mathematical modal logic: a view of its evolution* [88], Goldblatt wrote that Leibniz considered the concept of *possible worlds* and Wittgenstein spoke of *possible states of affairs*. He stated that, in [89], Saul Kripke provided a semantic model for a modal logic, which is often referred to as *possible world semantics* because it allows for different propositions to be true in different states, which are connected to each other by an accessibility relation. While this is not the only semantic for modal logics (cf., e.g., [90] for a non-Kripke semantics), this approach (to interpret modal formulae in models of possible worlds, with different state descriptions and in a relational structure) is exactly what we need (cf. Section 4.4), and thus the semantics we provide.

**Note:** Modal logic is consistent, complete and decidable [91]. Since ML extends PL, so is PL.

### 3.3. Game Theory and Rational Choice

Hastie and Dawes [77] dated the idea of rationality to the Renaissance and the middle of the 16th century. In the following century, Pascal and de Fermat devised an optimal strategy for betting (in games of chance) [92]. A more recent work on rational decision making—considered

to be a founding book for the field of *Game Theory* and already mentioned above (Section 2.2)—is “*Theory of Games and Economic Behaviour*” [42] where the authors compared economical behaviour to mathematical notions of strategic games [93]. With respect to game *theory*, Davis [93] agreed that “[t]he definition of game theory is very broad”. It is a collection of rigorous models aiming to explain interactive decision making [40] and a collection of analytical tools to understand the interaction between decision makers [94].

Amongst the decisions investigated by game theory is, e.g., strategic bargaining behaviour. There are many experiments from psychology that show humans are likely to punish *unfair* behaviour [95]. However, *fair* and *rational* are often conflicting behavioural stances. Without wanting to open this philosophical debate here, it should be noted that one underlying assumption of the term *game* (and in the field of game theory in general) is that an action is aiming to maximise the agent’s *pay-off* or reward. The fact that Osborne and Rubenstein’s model for rational decision [94] includes the individual player’s preference over the outcomes of actions directly implies that there is not one global preference shared by everyone (a generally adopted and supported understanding of what *fair* is). This means that their model for rational choice is concerned with rational, but not necessarily fair, decisions and that within this model the punishment of unfair behaviour can be irrational. Indeed, humans often fail to behave rationally [92]. In addition, Tversky and Kahneman [69] famously showed that the subjective evaluation of outcomes itself can be inconsistent when its outcome depends on a reference frame: differently framed versions of the same mathematical choice can illicit opposing outcomes [96].

Furthermore, the most rational decision, as calculated in a mathematical model such as the ones provided by game theory, is not always the most beneficial in the real world: for example, in October 1962, the *Cuban missile crisis* brought the world to the brink of nuclear war. Because of the magnitude of the payoffs (i.e., the potential consequences of the decisions taken, and the implications thereof), one would expect the decision making process to be guided by rational considerations. Indeed, “[t]he Cuban missile crisis is often held up as a model of rational decision-making” [97]. However, recent analysis has indicated that the decision making process that led Kennedy to the actions which eventually de-escalated the crisis were not rational in the traditional sense; only small differences in the circumstances would have led to very different events, potentially to global nuclear war [97]. Another example is the fact that, in the 1950s, von Neumann, using game theoretic models, argued for a pre-emptive and unprovoked nuclear first strike against Russia, something he saw as the most beneficial move in the *game* of nuclear proliferation [54]. Since then, the view of the underlying model has changed: Sagan Sagan [98] claimed that “*preventive war [is un]likely to lead to a safer nuclear future. Given the gravity of the risks we face, careful and steady movement towards global nuclear disarmament should be our goal*”.

These examples shows us that we can only use a theory to discuss its underlying model, whether the inferences made are then transferable to the real world depends on how adequate the model is. This is a very important consideration: we use common models and assume the existence of, e.g., preferences, but make no claim about these models and preferences being representative for the real world and for how humans really think, reason and, ultimately, behave. The latter is left to the experts; our stated aim is to assist them by providing a formalism, a *tool* so to speak, for their investigations.

### 3.3.1. Game Theory

Game theory is concerned with rational decision making. Its founding book [42] was written by a mathematician and an economist, which indicates what the field is about: mathematical models for social interactions which are on the one hand very well defined and on the other driven by a preference which is easily understood and modelled: the direct benefit of the acting entities.

Game theory is a decision making tool for scenarios where the individual choice and an element of chance are not the only deciding factors for the outcome of actions: the actions of others or changes in the environment (no matter how predictable) are included in the model [93].

Its formalism is mathematical, but the concepts and ideas modelled in Game theory are not inherently mathematical [94]. The descriptions and definitions of Game theory are stated formally to avoid ambiguity. The concepts that are captured can be understood intuitively (cf. [93]). Because of the interactive nature of the sequences of decisions we consider, we call the situations which we investigate *games*.

In these games, we have *players*, i.e., decision makers, whose actions and decisions are presumable guided by a *strategy*, i.e., a plan that allows the player to choose an action in all possible circumstances that can arise. Throughout the game, and especially at the end, there are *payoffs*, i.e., a reward or a punishment for each player, which can be compared between players to determine a *winner*.

Osborne and Rubenstein [94] provided a more formal model for rational behaviour (in the context of games): a set  $\mathcal{A}$  of actions has a set  $\mathcal{C}$  of consequences, there are rules  $\mathcal{R}$  which define the consequences of actions (or, as we below (cf. Definition 6 in Section 4.4.1) call it, determine transitions between the individual states of the game) and some preference function  $\succ$  orders these consequences according to how desirable they are. The latter is often determined by a *utility* measure. All the above form the basis for the formalisms presented in Section 4.3 as well as Definitions 13 and 14 in Section 5.6.

### 3.3.2. Utility Theory

Von Neumann and Morgenstern [42] acknowledged the “*conceptual and practical difficulties of the notion of utility*”. Our use of utilities permits us to use a straightforward numerical representation; a full philosophical discourse is outside the scope of this work, as it was outside the scope of [42].

Suffice it to say that, as marginally mentioned Osborne and Rubenstein [94], preferences are sometimes (as in our work) defined in terms of a numerical value given to the individual consequences of actions, expressing the *utility* value of this consequence. A preference function is then simply a means to calculate for any two consequences which, if any, is preferable to the other.

Let us briefly discuss what we understand by the term *utility value*: as Davis [93] pointed out, before making a decision to *get what we want*, we first have to have a clear understanding of *what it is* that we want. There is the philosophical argument that this is ultimately not the subject of a conscious decision [99], along the lines of: *we are free to do that which pleases us the most, but we are not free to decide what it is that pleases us*. Moreover, we change our mind when the context (but not the choice) changes [100]. Experiments have shown that humans are neither consistent in *what pleases them* [101], nor are they pursuing their desires consistently [102]. These are considerations which are not relevant for our use of utility: our aim is to provide a simple model for behavioural artificial intelligence (cf. Section 5.6). Whether this model meets the standards of psychologists is not relevant, as we argue that the model is such that it can be amended by practitioners to meet their standards. For us, it suffices to adopt the following basic assumptions from [103,104]: with respect to assigning a utility value to the consequences of actions, all outcomes are comparable and for any two outcomes, the preference is either clearly for one over the other, or for both (such that both outcomes are equally preferred and thus equally acceptable). Furthermore, preference and indifference are transitive (if  $A$  is preferred over  $B$ , and  $B$  is preferred over  $C$  then  $A$  is also preferred over  $C$ ). Since we are using a mathematical formalism already, the domain of numbers lends itself to express utility. By assigning a finite number of numerical values to any outcome, we add nothing new but facilitate the representation of a preference as an already intuitively understood operation (*equals* or *greater than*) [93].

### 3.3.3. Subjective Expected Utility (SEU) Theory

In the 17th century, it was proposed that, to calculate the risk—and prevent the worst—of harm, two subjective questions more than any other are considered: how much harm would this cause and how likely is it that this is going to happen (to me) [92]. There are a number of issues with simply assigning a utility value to consequences. As extensively discussed in the literature (cf. [102]), one of the issues is that there is often no certainty for an outcome to result from a specific action, only a probability. Another is that the evaluation of an outcome can strongly depend on the context [105].

Simon's model [16] of rational choice does allow probabilities for outcomes but assumes a constant utility for them. This model is briefly discussed here because it will be used in Section 5.6 to model rational behaviour. Simon himself pointed out that there are two aspects of any research on rationality: capturing what should rationally happen (normative) and what actually happens (descriptive) [14].

In [16], Simon suggests that a model for rational behaviour requires some of these elements:

- A set of behaviours from which to choose
- A subset thereof to consider
- A set of possible futures, resulting from the different behaviours/actions
- A payoff function representing the utility
- A function to determine the outcome a certain action will bring about
- Some information about the probability of that particular outcome occurring

Simon's model is very similar to our formalisms presented in Definition 13 because they were inspired by it. Generally speaking, it has to be understood that SEU is just a theory, that there are a number of theories (there is a good summary of them in [106]) and that there is criticism to SEU [107], notably, the work on *Prospect Theory* by Kahneman and Tversky [61]. Furthermore, prominent scholars such as Gigerenzer and Goldstein [36] have argued that there is no single model for human decision making but that we use a plethora of approximation techniques and heuristics to navigate every day life. We do not pass judgement on the appropriateness of SEU over other theories; it is used for this work because it captures aspects of interest to us but not more. Distinguished researchers have spent their careers on designing models for human behaviour and we are not trying to compete with them.

#### 4. Games as a Discreet Environment for Controlled Behaviour Assessment

In the previous section, we have introduced models for behaviour and provided a formalism to describe specific behaviours. We have also discussed a certain behaviour type as being prominently discussed in the literature. In this section, we discuss computer games.

- (a) We open by mentioning the use of games in psychology and education in general.
- (b) We introduce a certain type of game which fits our purpose. The aim is to advocate the use of games as controllable environments where test subjects can be subjected to comparable decisions (with the obvious aim to then record and compare these choices).
- (c) The bulk of this section formally defines what we understand to be a game in Section 4.3.
- (d) The bulk of this section provides a formal model of games in Section 4.4. This model can then be used to interpret behavioural statements (which is defined in Section 5) when expressed in the formalism presented in Section 3.

##### 4.1. Psychology and Computer Games

The behavioural activity of engaging in *play* is considered by, e.g., Brown [108], to be a fundamental basis for development in complex animals, on par with the act of sleeping and dreaming. According to Bruce [109], it can be a significant part of maturing from children to fully rounded adults and a strongly determining factor in the shaping of a functioning member of society. The literature lists many positive aspects of the use of games for serious purposes. Although they are often marginalised as such, games have never been *just* a children's medium [29]. Games have been used to great success to train complex problem management [110]—and problem solving [111]—abilities as well as practical and reasoning skills [112]. When used appropriately, they can significantly reduce training time and demands on the instructor [113]. Since games are generally something many people appreciate [114], they can have the advantage of maintaining high motivation levels in the learners [115]. The act of rehearsing is inherent to many games and as such is experienced as a pleasant repetition and not as boring rehearsals or automaticity training. For this reason, games have been widely used for decades now by large firms and companies to train their employees. Games have

been used in many areas as training and simulation tools: military training, teaching exact sciences (specifically mathematics), training in software engineering, computer science and information systems as well as medicine ([116–124], respectively).

We refer to [115] for a detailed account of the importance of intrinsic motivation to the designer of serious (computer) games. Two comparative studies, conducted in 2005 and 2007 have shown that challenge, curiosity and cooperation consistently emerged as the most important motivations for playing computer games, suggesting that appropriately designed games have a large potential to be suitable evaluation tools. For the full report on these findings, we refer the reader to [125,126].

The term *serious games* is not confined to education [127]: so-called *business games* were already proposed for research in the 1960s (e.g., [128]) and 1970s (e.g., [129]). It is safe to say that these games have been analysed from many different perspectives, both negative (e.g., aggression, violence or gender stereotyping) and positive (e.g., skills development, engagement or motivation) [130].

#### 4.2. Resource-Management Games for Serious Games

Resource-management games (RMG) are games in which a player is in charge of some coordinated effort in some simulated world. Perhaps the most famous example of this genre is SimCity, which has enjoyed so much attention that it is even used for urban planning and simulation [131]. Recently this type of game is being developed both as entertainment as well as a tool in professional contexts [132]. The overall goal of resource-management games is to maximise the outcome of the coordinated effort. In these games, the player often has to reach a number of intermediate goals and, to do so, the player has a limited number of choices to attain these sub-goals while using limited resources. This commonly forces some sort of tradeoff which requires the player to plan ahead for future actions. Therefore, such games provide a suitable platform to create controlled environments to study the interaction between the participating entities.

Malone [115] provided a detailed account of important aspects of intrinsic motivation in the design of serious games. It suggests that intrinsic motivation is created by four individual factors: challenge, fantasy, curiosity and control; as well as three interpersonal factors: cooperation, competition and recognition. Interestingly, these factors also describe what makes a good game, irrespective of its educational qualities. This parallel between what makes a good gaming experience and a good learning experience is also identified by Gee [133]. We are confident that RMGs can be designed to meet these requirements and are therefore an ideal type of game for our needs because:

- *They are challenging* due to restricted or limited resources, location and time, the need to plan ahead and the multitude of potentially conflicting objectives.
- *They stimulate* the fantasy by putting the player into an unfamiliar and imaginary position.
- They constantly require the player to control issues arising from the continuity of the game and from the actions of competing AI or human players. Choosing a behavioural strategy in response to actions of others is a substantial part of the game.

The two games presented in Section 6 as proof-of-concept implementation for either a formalism to capture (and evaluate in an automated manner) player behaviour (cf. Section 6.2) or for a rudimentary implementation of a behavioural game-playing AI (cf. Section 6.3) are both resource-management games.

#### 4.3. Defining a Game

Since 2005, there has been the *General Game Playing Competition* (<http://games.stanford.edu/>) (GGPC) ([134]) where programs compete in playing games. The games to be played are unknown in advance and the programs compete different types of games expressed in a Game Description Language ([135]). Other approaches and frameworks exist but there is no unanimously accepted definition for the concept of games in the literature, and some argue such a definition cannot exist. The definitions provided in this section are carefully chosen to suit the intended scope and domain of the presented formalism. They are not meant to cover any and all aspects of games. In the context



of this article, a game is played by a finite number of entities engaging in it, which we call *players*. There are no restrictions on their number except that we consider only finite groups of players and we require a minimum of one player for any game.

Here, we first discuss how to capture the essence of a game. This is not meant to be an algorithmic way how to actually compile a complete description of a game but is intended to serve as a definition of concepts. In the broadest sense, any game of the type we are considering can be defined by the set of all possible ways it can be played. The transition (by legal moves) from the start configuration of a game to an end state is what we call a *history*. Such a history can be seen as a sequence of admissible *states* the game can be in. To maintain the abstract nature of our approach, these notions are kept general for the moment. The relevant definitions are provided as they are needed.

The transition from one state to another happens after some action is taken or some event takes place. Since we already have *actions* as a term used in the context of the TACT paradigm (cf. Section 3.1), we call these transitions between states a *move* instead. In a game, there are traditionally *rules* determining the available moves, which are formalised in this article as a function mapping a state and an action to another state (see Definition 3).

**Definition 1** (Players, states and moves). *A game  $\mathcal{G}$  consists of a set  $\mathcal{P}_{\mathcal{G}}$  of  $n$  players, a set of states  $\mathcal{S}_{\mathcal{G}}$  (containing a starting state  $s_{start}$ ) and a set  $\mathcal{M}_{\mathcal{G}}$  of  $m$  move functions  $\mathcal{M}_i : \mathcal{S}_{\mathcal{G}} \rightarrow \mathcal{S}_{\mathcal{G}}$ . That is,  $\mathcal{M}_i(s)$  is a set of states to which a move from  $s$  is allowed. States  $s$  such that  $\mathcal{M}_i(s) = \emptyset$  are called “final states”.*

In our definition, we also include the move the player made, i.e., the transition from one state to the next. While this is not of direct use to us in the implementations presented in Section 6, it is included here to maintain the general nature of the approach.

Above, we define both the set of states as well as the set of possible moves to be finite. A history of a game however could be infinite if we allow the return to previously visited states, i.e., if we allow cyclic games. For the purpose we have in mind, infinite games are not relevant. Either way, as defined though the complete set of admissible histories (and as a logical consequence from the finiteness of the set  $\mathcal{S}$ ), there is, for any state, a finite subset of  $\mathcal{S}$  which can be directly reached from it. This accessibility is formally expressed by the relation  $\mathcal{R}$ , defined in Definition 3.

**Definition 2** ( $\mathcal{H}$ , a set of histories). *Let  $\mathcal{H}_{\mathcal{G}} = \{h_1, \dots, h_{n_{\mathcal{H}_{\mathcal{G}}}}\}$  be the set of admissible histories of game  $\mathcal{G}$ . An arbitrary history is a sequence of transitions  $\langle \text{state}, \text{move}, \text{state}_{next} \rangle$  between states:  $h_i = (\langle s_{start_i}, m_{0_i}, s_{1_i} \rangle, \langle s_{1_i}, m_{1_i}, s_{2_i} \rangle, \dots, \langle s_{n_i}, m_{n_i}, s_{end_i} \rangle)$  with  $s_{start_i}$  a starting state of the game and  $s_{end_i}$  a final state.*

Note that by this we assume that the information contained in the description of a state suffices to fully define it. That is, if a state is reachable from a given state in one incarnation of the game, it is always an admissible next state from there. This is because the rules of the game do not change, and thus the relation between states does not change either. This will be of importance in Section 5.6 as we consider the rational process of deciding on a suitable next state based exclusively on the information contained in states (and not based on the history or other parameters of the game).

At any stage throughout a game, there is one (partial) history that has led to the current state (the past which has already happened) but there may be a series of branching possible histories (the future which has yet to be played). This is of interest as we use our model for games in two ways: firstly, we evaluate the behaviour of players throughout the game (cf. Section 6.2), in which case we consider the single history of the game, as it was played. However, for the behavioural AI playing the game (cf. Section 6.3), the consideration of the possible future histories is relevant for the process of making a rational or behavioural decision on how to play.

From the above, that is, from histories, states and moves, we can derive the rules that govern the game. However, contrary to the intuitive understanding of rules in a game, we do not attempt to formalise the set of instructions that normally constitutes the set of “rules for a game”. Instead, by defining, for any existing state in a game, all reachable next states and the corresponding action,

we effectively cover all relevant rules of the game. Our choices are motivated by the formalism proposed in Section 5 which is based on modal logic (cf. [85] and Section 3.2.2) and game theory (cf. [94] and Section 3.3.1).

By locating all occurrences of a state  $s_i$  in all histories in  $\mathcal{H}$ , we can collect the set of states reachable as next state from  $s_i$ . By taking the actions under consideration, we can partition such a set into proper subsets that represent the states reachable from  $s_i$  by performing a specific action  $a_j$ .

We do not mean to imply that all games included in these histories have to be played before they can be included. The intention behind the concept of a history is that this is a concept, not the list of all games that have been played to this stage. There is such a set, and from that set we could derive the rules. The approach is along the lines of the game theoretic evaluation of, e.g., the game chess, where nobody is expecting that the researchers actually have a full list of all possible chess games.

**Definition 3** (The rules of a game  $\mathcal{R}$ ). Let  $\mathcal{R}_G = \{\mathcal{R}_1, \dots, \mathcal{R}_n\}$  be the set (of sets) of rules, one for each player  $i \in \mathcal{P}$ . Each of these sets  $\mathcal{R}_i$  is itself a set, containing elements of the form  $(m_j, s_p, s_q)$  (i.e., a move and the state in which it is performed as well as the resulting state), on which we impose the restriction that  $s_p \neq s_q$ .

This gives us for each player a full list of all state-transitions this player can bring about. The distinction between players and moves is relevant as players may assume different roles in the game and thus not be allowed to make the same types of move, however, their moves may overlap. Note that the above given definition for the rules of a game does not allow for moves that result in no change of the state of the game (reflexive relations), which is (in the context of this article) acceptable because we focus on the active behaviour within the game. Unless remaining passive is considered a conscious action, any action (behavioural decision) should have a consequence. This is a design choice, and if such reflexive relations are wanted, they can be used, and the model does not prohibit this. Amending Definition 3 to that effect has no impact on the overall formalism.

We can now define the concept of a game. Since we only consider states that can actually be reached (i.e., occur in the history) and moves which are actually performed at least once in at least one history, the set  $\mathcal{H}$  does define both the sets  $\mathcal{S}$  and  $\mathcal{M}$ . Furthermore, since we constructed the set of rules  $\mathcal{R}$  from  $\mathcal{H}$ , this is also indirectly included in our definition of a game:

**Definition 4** (A game  $\mathcal{G}$ ). A game is defined by the tuple  $\langle \mathcal{P}, \mathcal{H}, p \rangle$  with:

- $\mathcal{P}$  the set of players  $\{1, \dots, n_{\mathcal{P}_G}\}$ .
- $\mathcal{H}$  the exhaustive set of histories  $\{h_1, \dots, h_{\mathcal{P}_G}\}$ .
- $p$  a function  $p : (\mathcal{H}) \rightarrow (\mathbb{R} \times \dots \times \mathbb{R})_n$  mapping histories to individual payoffs for the  $n$  players.

The only thing new here is the *payoff function*. This function will assign a value to the victory of the individual players. This can range from simply indicating *win* or *loss* (boolean 1 or 0) to a function as complex as the game requires. As explained above, we can derive  $\mathcal{S}$ ,  $\mathcal{M}$  and  $\mathcal{R}$  from  $\mathcal{H}$ .

We do not include time constraints (e.g., maximum time for a move, etc) in our considerations, as we do not need them for the application we have in mind. Here, we first provide a model for these games, and then discuss the formalism that allows us to describe the players' behaviour.

#### 4.4. Modelling a Game

We model games on relational structures, often called *Kripke* models which are explained in Section 4.4.1 [136]. These are often used to provide semantics for modal logic (Section 3.2.2) and are very similar to models used in game theory (Section 3.3.1), both of which are used here. We briefly discuss disjoint sub-models (Section 4.4.3), specific instances of games (Section 4.4.5) as well as games with uncertainties (Section 4.4.6).

#### 4.4.1. Possible World Models (Kripke Semantics)

We formally introduce  $\Phi$  (which we have already identified before, cf. Section 3.2.1) the collection of all variables that fully describe any individual state of the game:

**Definition 5** (All relevant aspects of a game  $\Phi$ ). *For any state of a game, there is a finite set of atomic statements/variables capturing all its relevant aspects, and we define  $\Phi$  as the smallest set of all such statements.*

Regarding individual elements in  $\Phi$ , we require that they are *atomic*, i.e., not the combination of smaller statements. The statement “it is raining and it is Sunday”, e.g., could be the combination of two atomic statements “it is raining” and “it is Sunday”, in which case it would not be in  $\Phi$ .

To provide a formal model for games, we only represent the underlying structural aspects of a game. We argue that it suffices to represent, for any state of the game, which other states are reachable from it (cf. Section 3.2.2 where possible world semantics are discussed). This abstract modelling suffices for the analysis of games in which, e.g., the number of options available to the opponent is of importance. This could be because rendering the opponent without possible moves constitutes a win or, more practically thinking, because we might not have the resources to compute the full range of options so that restricting the opponent’s moves as much as possible may give us a better reasoning position.

Given a game  $\mathcal{G}$  we now define a model  $\mathfrak{M}$  for that game as follows: Recall that  $\mathcal{G}$  is defined as a tuple consisting of the set of  $n$  players,  $\mathcal{P}$ , the set of histories  $\mathcal{H}$  (sequences constructed over the set  $\mathcal{S}$  of states and moves) and a function assigning payoffs to the players with respect to a history.

We use standard concepts and definitions from modal logic (cf. [85] and Section 3.2.2) to capture the relational structure of a game. Consider that we are not interested in all the details of the specific game being played, but only the course a game can take. We start by introducing the idea of a so-called *frame* [85], i.e., a formal structure of a game without the information  $\phi$  [85] assigned to the individual states. This information can be added later, transforming such a frame into a (less generic) model of a specific game. The reason for considering the formal structure of a game is to get a better understanding of the game. This enables the game designer to verify whether a game can become, e.g., cyclic, or, more importantly, whether a game can be seen as a multitude of identical sub-games (structure wise), where, e.g., only the order of the players differs. Through this, the representation of the game can be reduced, which would in turn be expected to reduce the computational demands. The argument is that a *frame* captures all possible incarnations of a game that follow the rules, applying automated reasoning to this structure can vastly reduce the size of the search space. We do not insist on this model, but our section on game design (Section 6.3.3) motivates our use of frames.

We are interested in the possible successor states for a given state. To this end, we rewrite the information contained in  $\mathcal{H}$  to  $\mathcal{S}$  and  $\mathcal{R}$  and define the formal structure of a game over just these two:

**Definition 6** (Formal game structure  $\mathfrak{F}$ ). *Let  $\mathfrak{F}_{\mathcal{G}}$  be the formal structure for game  $\mathcal{G}$  with  $\mathcal{G} = \langle \mathcal{P}, \mathcal{H}, p \rangle$ . As such,  $\mathfrak{F}_{\mathcal{G}}$  is called a frame of game  $\mathcal{G}$  and defined as the tuple  $\langle \mathcal{S}, \mathcal{R} \rangle$  with:*

- $\mathcal{S}$  the set of states  $\{s_1, \dots, s_i\}$ .
- $\mathcal{R}$  a set of rules  $\{\mathcal{R}_1, \dots, \mathcal{R}_{n_{\mathfrak{M}_{\mathcal{G}}}}\}$ .

Analogous to the terminology used in [85,137] (cf. Section 3.2.2), we include a *valuation*  $\mathcal{V}$  that bijectively maps state names to subsets of  $\Phi$  (defined above, Definition 5), i.e., some set of statements  $\Phi_j$  (with  $\Phi_j \subseteq \Phi$  and such that  $\Phi_j$  is a complete description of state  $s_j$ ).

**Definition 7** (A valuation  $\mathcal{V}$ ). *We introduce a valuation  $\mathcal{V}$  to assign a subset of  $\Phi$  to each state of a game:  $\mathcal{V} = \{(s_1, \Phi_1), \dots, (s_i, \Phi_i)\}$  with  $\Phi_1, \dots, \Phi_i \subseteq \Phi$ .*

Valuations are not related to payoffs and other game related concepts. They are the assignment of truth values to propositions. Using frames and the valuation, we introduce the notion of a model:

**Definition 8** (A model  $\mathfrak{M}$ ). We introduce  $\mathfrak{M}_G$ , a model of  $G$ , defined by  $\mathfrak{F}_G$  and a valuation, i.e., the tuple  $\langle \mathfrak{F}_G, \mathcal{V} \rangle$  or  $\langle \mathcal{S}, \mathcal{R}, \mathcal{V} \rangle$  with :

- $\mathcal{S}$  the set of states  $\{s_1, \dots, s_i\}$ .
- $\mathcal{R}$  a set of rules  $\{\mathcal{R}_1, \dots, \mathcal{R}_{n_{\mathfrak{M}_G}}\}$ .
- $\mathcal{V}$  a valuation  $= \{(s_1, \Phi_1), \dots, (s_i, \Phi_i)\}$

#### Practical Considerations

Since we construct the set states of the game from the histories of a game, we consider only those states of a game that can actually occur. In chess, for example, there are board configurations that, while consisting of legal placement of figures of the game, can never actually occur in a game. In practice, however, one would most likely not enumerate all possible histories but instead simply create the superset of all states and then, through some algorithm, exclude those that are unreachable from any starting position. By this, we do not mean that this set would actually be created in the sense that it would contain all such states. Instead, we would define how the elements of this set would look like and then place some restrictions on them. To use the example of chess again, this would be a description of the placement of pieces on the board, restricted by some constraints.

#### 4.4.2. Potentially Infinite Games

We already mentioned above that our definition of a model covers games with potentially infinite histories since the model can be cyclic. In fact, all we require is that  $\Phi$  is finite, i.e., that the number of atomic facts we are willing to consider in our model is finite. From that, the finiteness of  $\mathcal{S}$  follows (consider that the maximum number of game states in which such facts are either true or false is limited (in fact, there are at most  $|\Phi|^2$ ) if we exclude identical ones). That in turn makes  $\mathcal{R}$  (see Definition 6) finite, while allowing histories of infinite length as players can revisit previous states.

Formally, the property of being cyclic is easily characterized. The existence of such histories might be an important factor since they may constitute a non-losing strategy, i.e., an approach that enables a player to prevent defeat. As an example consider the rule in chess that states that after a certain number of repetitions the game ends in a draw. For game designers, it should be of interest to verify that no such non-losing strategy exists. Note that the existence of a cycle does not necessarily imply that a player can force the game into this cycle.

#### 4.4.3. Disjoint Submodels

As stated earlier, we initially defined a game as the set of all admissible histories. For a game such as chess, this means that we include all rule abiding sequences of moves from the one starting position of the game. However, for other types of games, this might mean that we include the histories for *all* possible incarnations of the game, and thus far more than any one actual game might have. To illustrate this, consider the game of five-card poker where the reader might be dealt four aces or four kings, yet never both of these quartets in a single game. However, our definition of a model covers both. Figure 2 illustrates this: this model for poker consists of a number of (very similar) submodels which have no connection between them. In such a case, we can restrict our considerations and game design efforts to the submodel, which may considerably reduce the required work. Bear in mind that we are interested in games which we designed and where the existence of submodels may be a design choice (cf. the second proof-of-concept game *SoxWars*, discussed in Section 6.3.3).

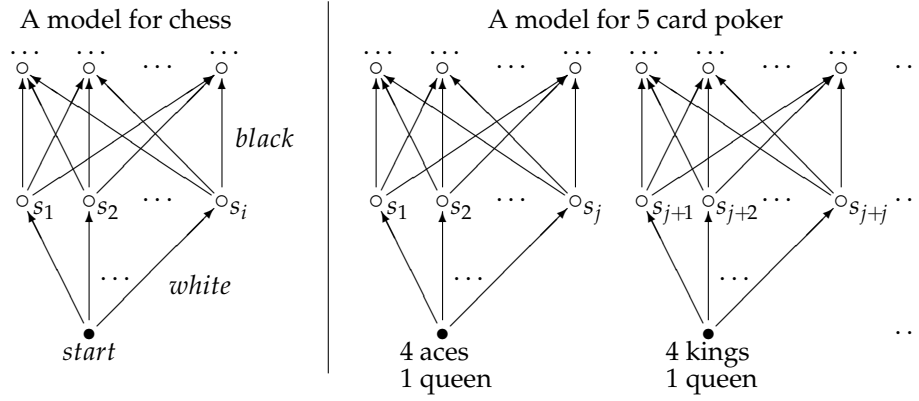


Figure 2. Two heavily simplified models for chess and poker.

Depending on the game, sometimes a number of facts can effectively *characterise* the submodels (e.g., hands in poker), meaning that some properties of the model hold only in one of its submodels. Such a defining characterization may arise from a combination of facts which do not have any obvious relation with one another. We refer the reader to the practical aim of this work: we use games as means to an end, namely to create carefully controlled environments within which we want to define and control behaviour. Therefore, the choice of which behaviours to focus on, and the decision on how to structure the game around this will be subject to change. This intended use of the formalism suggests that the game itself will be designed carefully, and may be designed to match certain properties or allow the characterization of submodels. We discuss submodels here because they can provide a computational advantage over complete models: for example, in Definition 13, we use the notion of *consequences* and nothing prevents us from considering reaching a specific state (or a set of states) in a submodel as a desirable consequence. This will effectively allow us to apply the decision making process shown discussed in Section 5.6 to a submodel instead of the entire model.

We define submodels intuitively as follows: a submodel is a model that does not share any of its states with the rest of the model (with respect to accessibility, i.e., the worlds in the submodel are not reachable from any world that is in the model but not in the submodel). The submodels of  $\mathfrak{M}$  will thus partition  $\mathcal{S}$  and  $\mathcal{R}$  (and, since the valuation is on the elements of  $\mathcal{S}$ ,  $\mathcal{V}$ ) such that each world in a submodel is accessible exclusively from worlds within that submodel:

**Definition 9** (A set of subframes  $\mathfrak{F}^{sub}$ ). Let  $\mathfrak{F}^{sub} = \{\mathfrak{F}_{sub_1}, \dots, \mathfrak{F}_{sub_n}\}$  be the set of subframes of  $\mathfrak{F} = \langle \mathcal{S}, \mathcal{R} \rangle$ , then we require that  $\mathcal{S}$  is partitioned by the mutually exclusive sets  $\mathcal{S}_{sub_1}, \dots, \mathcal{S}_{sub_n}$  of the subframes and that all rules  $\mathcal{R}_1, \dots, \mathcal{R}_i$  in  $\mathcal{R}$  are partitioned likewise by the corresponding sets in  $\mathcal{R}_{sub_1}, \dots, \mathcal{R}_{sub_n}$ . As we intend to define “disjoint” subframes, we also require that any such set  $\mathcal{R}_{j-sub_k}$  may only range over  $\mathcal{S}_{sub_k}$ .

Disregarding the order, there is exactly one such set of subframes for any frame (see Definition 6) and it is clearly the last requirement that captures our intention behind this partition as it is the one responsible for the disjoint nature of our frames.

The definition of the submodels constructed over these frames is therefore:

**Definition 10** (A set of submodels  $\mathfrak{M}^{sub}$ ). For model  $\mathfrak{M}_G = \langle \mathfrak{F}, \mathcal{V} \rangle$ , the set  $\mathfrak{M}_G^{sub}$  of submodels is constructed by pairing the elements of  $\mathfrak{F}^{sub}$  with (mutually exclusive) partitions (as determined by the correspondence to the respective sets of worlds) of  $\mathcal{V}$ , i.e.,  $\mathfrak{M}_G^{sub} = \{\langle \mathfrak{F}_{sub_1}, \mathcal{V}_{sub_1} \rangle, \dots, \langle \mathfrak{F}_{sub_n}, \mathcal{V}_{sub_n} \rangle\}$ .

The reader will intuitively understand the requirement that each partition of  $\mathcal{V}$  should cover exactly those worlds included in the frame it is paired with. Since the sets of worlds in the subframes are disjoint (i.e., no world can be in more than one subframe), the partition of  $\mathcal{V}$  is mutually exclusive.



#### 4.4.4. Characterizing Submodels

Having defined the notion of submodels, we now have the ability to formalise which elements of  $\Phi$ , if any, characterize them. Above, we explain that the valuations of the submodels  $\langle \mathcal{V}_{sub_1}, \mathcal{V}_{sub_1} \rangle, \dots, \langle \mathcal{V}_{sub_n}, \mathcal{V}_{sub_n} \rangle$  of  $\mathcal{M}^{sub}$  are disjoint partitions of  $\mathcal{V}$ . The subsets of propositions assigned to the individual worlds are not under that constraint. Let us consider the collection of all propositions assigned to at least one world in a submodel  $\mathcal{M}_i$  and call it  $\Phi_{\mathcal{M}_i}$ .

#### Practical Considerations

We might be interested in similarities between individual incarnations of a game. As frames provide us with a purely structural view on a game, subframes allow us to compare instances of games on that basis. Figure 2 (right) sketches a model for poker. Now, as the poker-playing reader will agree, a game of poker always looks the same to an outsider, as the game does not offer a great choice of different moves. The factual difference in games lies in the individual cards held by the players. Were it not for the choice to either bid, call or fold, all games would look exactly the same for the outside observer, right up to the last action (revealing the cards). In our model, this results in a large number of identical subframes, all of them differing only in the distribution of the cards. One benefit of introducing subframes is thus that, while the model for poker is rather large, the frame we get after removing all identical subframes is so small that it can even be drawn on paper.

For example, for the proof-of-concept game *SoxWars* (Section 6.3.3), several variations on the sequence of actions in a turn were considered. To someone designing and implementing a game, it may be of great help to be able to represent a game at this abstraction level.

#### 4.4.5. Specific Instances of a Game

For games such as chess, the current model is already the model for any incarnation of that game. For most dice and card games, there are parts of the model that can never be reached in a game: for example, those states which differ from the hands that are actually dealt or the actual outcome of rolling the dice. In other words, in games involving the element of chance (such as poker), there will be disjoint submodels. To fix this, one could combine all submodels of a game under a master-root state, which may represent the state *before* the cards are dealt. However, since this does not fit in with the translation from  $\mathcal{H}$  to  $\mathcal{M}$ , we now introduce the model for an arbitrary instance of a game.

Above, we define  $\Phi$  as the set of all propositions needed to describe any history of a game. However, unless all games are actually the same, there will be some proper subset of  $\Phi$ , which we denote by  $\Phi'$ , that contains only those propositions that are relevant to the specific history in question. For example, the statement “*Player 4 wins the game*” is irrelevant for all cases of three players or fewer. While the claim of irrelevance might seem a bit strong, we point out that Definition 5 defines  $\Phi$  to be the *smallest* set of statements capturing all *relevant* aspects. Analogously, we do not require all propositions in  $\Phi'$  to hold everywhere in the model, but we consider them relevant, i.e., they could be true somewhere in some state in the (refined) model. This is, again, a consideration aiming to reduce the model of a game to the smallest required form. This is a relevant factor when implementing the game and as such, the argument and this section further advocates the feasibility of the suggested approach.

From a given model  $\mathcal{M}$  of a complete game, constructed over a set of propositions  $\Phi$ , we can construct  $\mathcal{M}_{\Phi'}$ , the model of a specific history (i.e., where  $\Phi'$  is the set of all true propositions).

**Definition 11** (Refining upon  $\mathcal{M}$ ). *Let  $\mathcal{M} = \langle \mathcal{S}, \{\mathcal{R}_1, \dots, \mathcal{R}_m\}, \mathcal{V} \rangle$  be a model (over the set of propositions  $\Phi$ ) and let  $\Phi'$  be the propositions we deem relevant ( $\Phi' \subseteq \Phi$ ). We construct the refined model  $\mathcal{M}_{\Phi'} = \langle \mathcal{S}_{\Phi'}, \mathcal{R}_{\Phi'}, \mathcal{V}_{\Phi'} \rangle$  as follows:*

1.  $\mathcal{S}_{\Phi'} = \{s | s \in \mathcal{S} \text{ and } (s, \Phi'') \in \mathcal{V} \text{ with } \Phi'' \subseteq \Phi'\}$
2.  $\forall \mathcal{R}_{i'} \in \mathcal{R}_{\Phi'} : (m_j, s_k, s_l) \in \mathcal{R}_{i'} \Leftrightarrow (m_j, s_k, s_l) \in \mathcal{R}_i \text{ and } s_k, s_l \in \mathcal{S}_{\Phi'}$
3.  $\forall (s_i, \Phi'_i) \in \mathcal{V}_{\Phi'} : (s_i, \Phi'_i) \in \mathcal{V}_{\Phi'} \Leftrightarrow (s_i, \Phi_i) \in \mathcal{V} \text{ and } \Phi'_i = \Phi_i \cap \Phi'$

Our notion for a specific instance of the game here is very loose as it depends on  $\Phi'$ . The above enables us to model the part of a game which we are interested in. This can be all incarnations of the game with at most three players or one specific game where Player 1 wins in Round 3, for example with a full house. Analogous to the above, we can also refine a model on the basis of insights gained during the course of a game or to reflect only a certain number of actions.

#### 4.4.6. Uncertainties

We distinguish two types of games with uncertainties: those during which the uncertainties are successively removed and those where the uncertainties are a constant part of the game. To illustrate this, we name Memory (in the game *Memory*, players take turns to publicly reveal two cards; failure to pick identical cards (a pair) results in the cards being turned over again with their location unchanged) as an example of the former and Perudo (in the game *Perudo*, players place increasing bids on the total sum of their individual hands) for the latter. Furthermore, there are games that fall under both categories (in the game *Texas Hold'em Poker*, players do not show their own cards until the end but throughout the game a certain number of cards relevant for the final outcome are openly revealed).

In all of these, we can describe uncertainty as follows: if a player  $i$  has imperfect information regarding a game, then  $i$  has not access to the whole valuation, i.e., there are states for which the player does not know the truth value of all the propositions. Let  $\mathcal{V}^i$  be the partial valuation available to player  $i$ . We can then construct model  $\mathfrak{M}^i$  as defined in Definition 8. However, because of this,  $\mathfrak{M}^i$  may now contain multiple states with identical valuation. If we collapse these onto each other (and update the relations appropriately), we get the subjective model for player  $i$ . If there exists an action in a state of  $\mathfrak{M}^i$  that (for player  $i$ ) leads to more than one possible state, we say that there is “uncertainty” for  $i$  in the game. This notion is rather important as such a subjective model may no longer be deterministic, i.e., a straightforward application of a greedy strategy (always maximizing the player’s payoff) may not be possible. This comes back in Section 5.6 where this is discussed in the context of AI players dealing with probabilities of outcomes (instead of deterministic ones).

**Definition 12** (A model for a subjective view). Let  $\mathfrak{M}^i$  be a model representing the knowledge of player  $i$  in a specific game of  $\mathfrak{M} = \langle \mathcal{S}, \mathcal{R}, \mathcal{V} \rangle$  as captured by valuation  $\mathcal{V}^i$  (which we require to meet the following constraint:  $(s_j, \Phi_j') \in \mathcal{V}^i, (s_j, \Phi_j) \in \mathcal{V} \Rightarrow \Phi_j' \subseteq \Phi_j$ ). From  $\mathcal{V}^i$ , we define  $\mathfrak{M}^i = \langle \mathcal{S}^i, \mathcal{R}^i, \mathcal{V}^i \rangle$  s.t.:

- $(s_j, \Phi_l) \in \mathcal{V}^i \Leftrightarrow (s_j, \Phi_l) \in \mathcal{V}^i, \forall (s_k, \Phi_l) \in \mathcal{V}^i, s_j \neq s_k \Rightarrow (s_k, \Phi_l) \notin \mathcal{V}^i$
- $s_j \in \mathcal{S}^i \Leftrightarrow (s_j, \Phi_l) \in \mathcal{V}^i$
- $\forall \mathcal{R}_i \in \mathcal{R}, \forall \mathcal{R}_i^i \in \mathcal{R}^i, \forall (a_j, s_k, s_l) \in \mathcal{R}_i$ :

$$\begin{aligned} s_k, s_l \in \mathcal{S}^i &\Rightarrow (a_j, s_k, s_l) \in \mathcal{R}_i^i \\ s_k \notin \mathcal{S}^i, (s_k, \Phi_l) \in \mathcal{V}^i, (s_m, \Phi_l) \in \mathcal{V}^i &\Rightarrow (a_j, s_m, s_l) \in \mathcal{R}_i^i \\ s_l \notin \mathcal{S}^i, (s_l, \Phi_l) \in \mathcal{V}^i, (s_m, \Phi_l) \in \mathcal{V}^i &\Rightarrow (a_j, s_k, s_m) \in \mathcal{R}_i^i \end{aligned}$$

Note that in this section we only provide a subjective model from a single player’s point of view. Extending this to cover all players while outside the scope of this work is certainly interesting and considered for future work. Accounting for what the human player can deduce from the information provided throughout the game is certainly of interest but would possibly pose a problem as well: If the AI players base their actions on what the human could know, the results can become dependent on whether the human player actually inferred this information. While this would be of great interest in the context of testing whether the human player does indeed make full use of his/her potential (or, indeed, testing the ability to use complex information), it is outside the scope of the work presented here. In this article, we do not address AI players that can reason about their knowledge regarding other players. However, the modal logic based formalisms are standard (cf. [137]) and including them is a straightforward extension to our proposed approach.

## 5. Formalizing and Evaluating Human and Machine Behaviour

It is our aim to propose an unambiguous and precise formalism to express a variety of aspects of behaviour in games. The description of *behaviour in a game* will be in terms of the aspects of *that game*. For example, whether the player is rolling the dice calmly or in an aggregated manner is *not* something we consider to be a behaviour in the game. The behaviours we *are* interested should all be expressed by the moves made in by the players in the game (e.g., that the player decided to move his queen instead of the king). Therefore, we start Section 5.1 by formally defining everything of relevance *in a game*. For this, we use the formal language of propositional logic which we introduced in Section 3.2.1. Once we have a formal language that allows us to make statements about a game, we consider (in Section 5.2) how to properly express behaviour (in a game) in this language. As there are a number of different models for human behaviour in the field of psychology (we discuss four in Section 2.4), we had to decide on one of these paradigms but it should be understood that it can easily be exchanged for another; the work presented here is intentionally kept open for personal preferences. The approach we chose to base our work on is the TACT paradigm (Section 3.1).

### 5.1. Formal Statements about Individual States of a Game

As discussed in Section 4 (see Definition 1), we construct our model for a game over a number of states (i.e., the different states the game can be in). Since we are considering computer games, we argue that any such game state will be represented in the computer by a finite number of variables, i.e., by the data in the computer or a subset thereof. To maintain the generic nature of our approach, we do not specify anything about these variables at this stage.

In computer programs, the atomic statements which we defined in Definition 5 are Boolean variables. Since we implemented this on a computer, we could ensure that any of these statements will, at any time, either be true or false. This is analogous to the properties of  $\Phi$  from propositional logic, as previously described in Section 3.2.1. To combine these propositions into more complex statements, we use the two operators “not” and “and” as defined in Section 3.2.1. As discussed in Section 3.2.1, these two suffice to express three other concepts which we introduce as abbreviations: “or”, “if . . . then” and “if, and only if”.

We thus have, in total, five connectors to build increasingly complex statements. While the “not” precedes statements, the other four connectors are placed between two statements, e.g., “it is raining” and “it is Sunday”. Note that we use “and” to enhance readability. We rephrase the *not* if necessary to enhance readability, e.g., *it is not true that “it is raining” and “it is Sunday”*. As we already see from this example, the quotation marks (“and”) may accumulate excessively when we construct longer statements.

Although at times convoluted, we can now express arbitrarily complex statements about individual states of a game. We can, furthermore, implement an algorithm (Section 5.3) to translate such formal statements into natural language statements and back. As shown for one proof-of-concept game (see Section 6.2), implementing this is straightforward and can be a great tool, amongst other things, to enable us to amend individual behaviour stances for an AI player, possibly even while the game is being played.

### 5.2. Formal Statements about Behaviour in Games

In Section 2.4, we discuss models of behaviour, and in accordance with one of these (namely, the ToPB, Section 2.4), we introduce the TACT paradigm (Section 3.1). In our work ([19–21,24]), we have repeatedly used this paradigm to implement formally defined descriptions of AI behaviour.

#### 5.2.1. Behaviour in General

As discussed in Section 3.1, Ajzen [73] proposed defining behaviour by identifying and distinguishing four aspects: Target, Action, Context and Time. While it might be difficult to identify all propositions that

could refer to actions (and, indeed, to separate all propositions according to the conceptual partitioning required by the TACT paradigm), in the context of computer games, actions should be the most straightforward of these four to identify; actions are directly related to the interaction between the game and the player. In many cases, we will simply be able to record the actions from the limited number of choices offered to the player. The target of the action will also be relatively easy to identify as it will be the object of the action, or the entity at which an action is directed. Both will be identifiable in the current state of the game. The remaining two (context and time) will be a bit more difficult: the context of the action or behaviour could be something that relates to the course of the game, the already played part of the game or it might even be related to potential actions or events in the future of the game. Deciding on those requires a specific game instance.

### 5.2.2. Behavioural Statements

Different states of the game can be distinguished on the basis of the different propositions that are true in them. The view taken is that there are certain statements about a state that constitute an action being performed by a player (e.g., knight takes pawn). Such an action will result in consequences (e.g., knight not at position a, knight at position b, pawn not at position b). We chose to model this such that the state where an action occurs is also the state where the consequence is true. The individual states differ thus both in the facts (consequences) that are true in them as well as in the actions that have brought about these consequences. There are a number of reasons for this choice. Here, it suffices to point out that in the model there is no temporal delay between the execution of an action and the manifestation of the resulting consequences. This is a simplification that can be removed: if this were desired such states could be separated into multiple states, which are then connected by an accessibility relation related to the actions. While this might be of interest to some investigations, it is not relevant for the work proposed here and thus the presented model does not include it.

The consequences may be determined by the target of the action. In the design of the game, there can therefore be a partitioning of  $\Phi$  such that there are two disjoint subsets  $\Phi_a$  and  $\Phi_t$  containing the respective propositions. This conceptual representation may not be intuitive. Our aim is to adhere to the TACT paradigm but the approach can be easily adapted to similar distinctions.

The game can be designed as a series of formulae of the form  $\phi_a \wedge \phi_t \rightarrow \phi_c$  ("if a certain action and a certain target, then a certain consequence"), i.e., by defining the impact actions and their targets have on the remaining propositions, collected in the subset  $\Phi_c$  (for context). The interpretation is that actions do not change other actions, nor do they change the target of actions. The changes that do happen have to be happening in terms of changes of truth values for the remaining propositions.

As mentioned before, a game is represented as a history, i.e., as a sequence of states. Any analysis of the behaviour of the individual players in this sequence will be undertaken from the viewpoint of the initial state, and after the game has been played. This facilitates the temporal aspect which the formalism is currently still missing: temporal statements about the behaviour in the game can only be made in terms of the evolution of the game, i.e., with respect to the sequence of states. Examples are *within the next five moves* and *at least once before the game ends*. As such, time is measured in states or, alternatively, using meta units which would have to be imposed on the model (e.g., rounds or years where the former is a time unit that is part of the game).

### 5.2.3. Nested Statements

It should furthermore be understood that formulae can be nested, i.e., that a formula  $\phi_t$  can contain not only a reference to the player that is the target of the action, it can also refer to previous behaviour of that player. This facilitates the formulation of statements such as *Player A playing Action 1 against a player that played Action 2 against Player A in the previous round*.

There are a number of issues with this, and we are making no claim towards this being a universal solution towards expressing behaviour in games. A large part of the specific implementation will depend on the specifics of the game for which it is designed, as well as the behaviours that are going

to be investigated. Any attempt to make the formalism applicable in the wider sense has resulted in a bloated list of definitions and conventions, of which only a few are ever relevant to a specific case.

#### 5.2.4. Complex Behavioural Statements

In addition to the individual behaviours, we can define classes of behaviour to describe complex behaviours though this is merely to enhance readability and intuitive understanding: consider that we have three behaviours  $\phi_1$ ,  $\phi_2$  and  $\phi_3$  which describe a specific behavioural stance (e.g., *playing cooperatively*). We can then introduce an abbreviation for them:  $\phi_{coop} = \phi_1 \vee \phi_2 \vee \phi_3$  and use it in behavioural statements.

This does not add to the language, but enhances readability and enables us to design complex behavioural statements. Especially with the aim of designing a behavioural AI, this may be very useful as it can facilitate the amending of high level behaviour (such as the mentioned *playing cooperatively*) with a few actions by the user. There is nothing that prevents this from being implemented to work even while a game is being played. The design and definition of such complex statements will be part of the overall design effort, and our work has shown that this will constitute a substantial amount of the theoretical work [138].

Examples of this, taken from the main prototype game (cf. Section 6.3) are:

- **Cooperative**
  1. Player  $i$  is not bidding against any player this round
  2. Player  $i$  is not bidding against any player that has played cooperatively the last round
  3. Player  $i$  is not bidding against any player (that has played competitively against any player that has played cooperatively the last five rounds) this round
- **Competitive**
  1. Player  $i$  is bidding against another player this round
  2. Player  $i$  is bidding against all other players the next three rounds
  3. Player  $i$  is only bidding against players (that have bid on Player  $i$  the last round) this round

The TACT labels in the above examples are:

**Action:** *Player  $i$  is bidding*

**Target:** *Player 1, Player 2, Player 3, Player 4*

**Context:** *Player  $j$  has played  $\phi$  (where  $\phi$  is another TACT statement)*

**Time:** *this round, last round, last five rounds, next three rounds*

These are simple examples, and there are obviously additional issues that have not been addressed here (implementation of quantifiers, consistency of statements, temporal aspects, etc. [138]). The presented approach is intentionally kept general.

#### 5.3. Automated Translation of Behaviour Statements

Two translations should be discussed here: Firstly, we briefly outline how the formalism can be translated, in an automated manner, into statements in natural language. This is to facilitate the use of the formalism as a tool for non-logicians. Secondly, a translation into a so-called *normal form* is given. The latter is used for the automated evaluation of formal statements in the context of a history and a state of the model if the game.

##### 5.3.1. Formalism $\Leftrightarrow$ Natural Language

The above introduced formal language PL has the advantage of being well defined and unambiguous. The translation to natural language is intuitively understood [139], cf. Table 3:



**Table 3.** Statements in propositional logic and the corresponding equivalent statements in natural language.

Connector	Usage	Operator	Usage	Rewritten
<i>not</i>	<i>"it is not the case that statement"</i>	$\neg$	$\neg\phi$	
<i>or</i>	<i>"statement1 or statement2"</i>	$\vee$	$\phi \vee \psi$	
<i>and</i>	<i>"statement1 and statement2"</i>	$\wedge$	$\phi \wedge \psi$	$\equiv \neg(\neg\phi \vee \neg\psi)$
<i>if ... then</i>	<i>"if statement1 then statement2"</i>	$\rightarrow$	$\phi \rightarrow \psi$	$\equiv \neg\phi \vee \psi$
<i>if and only if</i>	<i>"statement1 if (and only if) statement2"</i>	$\leftrightarrow$	$\phi \leftrightarrow \psi$	$\equiv (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$

We now briefly illustrate how to translate natural language statements into PL and back. The process is straightforward and we illustrate it here for completeness only:

Repeat the steps below until all natural language elements are removed.

1. Translate atomic statements into their propositions.
2. Replace occurrences of *"it is not the case that"* by  $\neg$ .
3. Rewrite *"statement1 and statement2"* to  $(\text{"statement1"} \wedge \text{"statement2"})$  and, correspondingly, *"statement1 or statement2"* to  $(\text{"statement1"} \vee \text{"statement2"})$ .
4. Sentences of the form *"if statement1 then statement2"* are replaced by  $(\text{"statement1"} \rightarrow \text{"statement2"})$ .
5. Finally, *"statement1 if and only if statement2"* is replaced by  $(\text{"statement1"} \leftrightarrow \text{"statement2"})$ .

The reverse is analogous to the above and thus omitted. If the use of brackets and quotation marks is implemented correctly, the process can be automated. Since the aim is to make the natural language output as readable as possible, the specifics regarding the use of brackets and quotes are left to the designer, who might, for example, prefer *"not statement1"* over *"it is not true that "statement1 holds"* or vice versa.

### 5.3.2. Formalism $\Leftrightarrow$ Normal Form

When it comes to the automated evaluation of behaviour statements, we eventually look at the individual propositions and use their truth values to determine whether a specific complex statement is true (or false) in its entirety. To do so, we rewrite the statements to contain only  $\neg$  and  $\vee$ . Alternatively (as explained in Section 5.4, we might want to rewrite any statement to contain only  $\neg$  and  $\wedge$ , but the process is analogous. The rewrite rules below provide semantically equivalent statements containing only the operators for *not* and *and*, listed in Table 3, above. This provides the algorithm to successively rewrite the operators.

In short, the following are the rewritten rules (for the normal form with  $\neg, \wedge$ ):

1.  $(p \leftrightarrow q)$  becomes  $\neg(p \wedge \neg q) \wedge \neg(\neg p \wedge q)$
2.  $(p \rightarrow q)$  becomes  $\neg(p \wedge \neg q)$
3.  $(p \vee q)$  becomes  $\neg(\neg p \wedge \neg q)$

while these are the rewritten rules (for the normal form with  $\neg, \vee$ ):

1.  $(p \leftrightarrow q)$  becomes  $\neg(\neg(\neg p \vee q) \vee \neg(p \vee \neg q))$
2.  $(p \rightarrow q)$  becomes  $\neg p \vee q$
3.  $(p \wedge q)$  becomes  $\neg(\neg p \vee \neg q)$

### 5.4. Automated Behaviour Evaluation

We now have a model for games as well as a formalism to express behaviour statements regarding the modelled games. To complete this section, we now discuss the means to evaluate such statements using the model as well as histories of games played in that model.

In Section 4.3, we define histories (Definition 2). These enable us to evaluate the behaviour of players by looking into the past, i.e., by considering a single history. In cases where we are looking

ahead at a number of options (i.e., at potential outcomes of future decisions), we can consider this as a series of evaluations (one for each considered option), each resulting in the already existing history and on last entry, namely the one from the considered next move. The evaluation of such a hypothetical history will be relevant for the design of behavioural AI (cf. Section 5.6) and will follow the same lines, and the computational cost to evaluate them will be linear in the number of choices considered.

As stated in Section 5.1, we consider every state of the game to be a subset of  $\Phi$  (the set of all atomic statements). In other words, we describe the current state of a game as a list of all those statements which hold in that state. Now, given the truth values for the individual propositions, we can then evaluate any statement, however complex, to either true or false. To do so, we first make use of the above sketched translation to rewrite any statement into a semantically equivalent one consisting only of propositions and the connectors  $\neg$  and either  $\vee$  or  $\wedge$ . Afterwards, we evaluate the new statement according to the truth values for the propositions it contains.

If we evaluate a statement that is enclosed by the *not* operator, we simply evaluate the statement and reverse the resulting truth value (true becomes false and vice versa). Upon considering the truth table for the  $\wedge$  connector and, e.g., the statement  $(p \wedge (q \wedge r))$ , we see that we can simplify this to  $(p \wedge q \wedge r)$ . The same holds for the  $\vee$  connector which also allows to omit brackets if only  $\vee$  connectors used). We can rewrite any statement into a series of sub-statements connected by the  $\wedge$  connector. These can be evaluated independently. The reason for this is that if we find a single sub-statement that evaluates to false, the whole statement will be false. Depending on the implementation, it might be the case that we expect most of the statements to evaluate to true and in that case it might be beneficial to rewrite the original statements into ones containing only  $\neg$  and  $\vee$ , because, if the whole statement translated into a series of sub-statements connected by  $\vee$  has a single sub-statement that evaluates to true, the whole statement will be. Either way, as the tests discussed in Section 6.2.4 indicate, the computational cost will be well within acceptable limit. However, the formalism could theoretically be used to describe very large numbers of rather complex behaviours in games of extensive length. Should the computational cost become an issue, the above considerations (i.e., which normal form to use) might be applied to reduce the time an automated evaluation will take.

### 5.5. Automated Generation of Consistent Behaviour Statements

Using the above, we briefly mention a useful tool which we can now build, using the formalisms and automated processes introduced so far: using a small number of *core* behaviours, we can generate new behavioural statements that are consistent with existing ones. Consider that we want to program the AI of the game to behave more individually, and that we thus would like to add a number of behaviour stances to the individual AI players. We can use such a module to generate statements which address a number of issues that are currently not addressed because we can automatically identify conflicts and contradicting statements.

### 5.6. Behavioural Artificial Intelligence

In the computer game industry, AI players are used not only to provide challenging opponents but increasingly to add to the playing experience. Formalizing and modelling behaviour is a time intensive process [140] but creating computer agents with human-like cognitive behaviour will greatly improve the reception of these agents by the human players [66].

Thus far, we have discussed rational and realistic behaviour and how to model either what a rational decision would be, or the decisions actually taken by intelligent beings: humans. In the context of this article, we would now like to broaden this view from human intelligence and behaviour to intelligent behaviour in general and to intelligence as exhibited by non-living entities such as machines and programs specifically. We argue that it is not required to have a complete model of the former to address the latter because “[w]e do not need to understand intelligence to create it” [141]. The model of rational choice proposed by Simon [16], the *subjective expected utility (SEU) theory* (cf. Section 3.3.3), has already been mentioned in Section 2.2. It has been considered when designing the presented

model for rational behaviour and there is an intentional notational and conceptual overlap between the presented model and Simon's work. Within a game, the players make their decisions based on their subjective evaluation of the consequences of a move (as opposed to considering only the state in the game which a specific move will bring about). As discussed in Section 2.1, this view is not undisputed in behavioural psychology. While some of the mentioned theories of behaviour have great importance as a theoretical framework and as a guideline to applications [142], we restrict our use to the SEU model.

While the actual state of a game is effectively the same for all players, they might have different goodness values assigned to it. The notions of a move and a decision are thus strongly connected but very different in nature. Below, we define the subjective functions that map actions to consequences, consequences to goodness values and finally consequences to utility values. In addition, we introduce a preference relation that allows players to rank consequences and thus actions in an order of preference.

Utility values are also used by Cowley [143] ("we can calculate a ranking for all the choices available to the player, based on the utilities associated with the game states produced by each choice.") and evaluating the goodness of a move has been calculated in game AI ever since complex games were considered.

When considering future actions, we will be considering possible histories (cf. Section 4.3, Definition 2), i.e., possible extensions to the one history that has already been played. This means that we will consider more than one *potential history*. Of course, when evaluating the already exhibited behaviour of other players (the decisions already taken), the game behaviour model has to consider the past, i.e., the already played part of the history.

When making a decision on the next move, i.e., when deciding which of the available options is the most preferable, we want the game behaviour model not only to consider the pay-off maximising consequences but also to act in accordance with its behavioural stance [144]. Rational behavioural decision making is modelled (similar to [16]) as the tuple  $\langle \mathcal{A}, \mathcal{C}, r, \sigma \rangle$ .

**Definition 13** (Rational decisions in games). *Let the model for rational decisions (moves) be the tuple  $\langle \mathcal{A}, \mathcal{C}, r, \sigma \rangle$  with:*

- $\mathcal{A}$  the set of actions  $\{a_1, \dots, a_i\}$  available to the player.
- $\mathcal{C}$  the set of consequences  $\{c_1, \dots, c_j\}$  of these actions.
- $r$  the rules, i.e., a function  $r : \mathcal{A} \rightarrow \mathcal{C}$  mapping actions to consequences.
- $\sigma$  a (subjective) strategy, i.e., a function  $\sigma : \mathcal{A}' \rightarrow \mathcal{A}$  mapping subsets of  $\mathcal{A}$  to an action  $a$  ( $\mathcal{A}' \subseteq \mathcal{A}$ ,  $a \in \mathcal{A}'$ ).

*If required, we indicate the respective player  $k$  in the subscript ( $\langle \mathcal{A}_k, \mathcal{C}_k, r_k, \sigma_k \rangle$ ). However, this is included for completeness sake only as players share  $\mathcal{A}, \mathcal{C}$  and  $r$  in most games. We omit this wherever possible.*

The simplification we make here is that knowledge about the actions, rules and their consequences is shared, that is that all players have the same information available to them. The argument for this is that, at this stage, it is unlikely to assume that the human players will maintain mental models of the different knowledge their opponents have; and even if they did that, it would increase the complexity of the exhibited behaviours far above the level this approach is investigating.

With respect to the set of consequences, we discussed uncertainties in a game in Section 4.4.6. These can be due to hidden actions by the other players or to unpredictable events such as rolling a dice. The presented formalism is capable of handling these extensions; however, including this in the model here will not add anything but make the description considerably more complex.

Therefore, it is in the strategy  $\sigma$  that the individual player might differ from the other players. Generally speaking, any action  $a$  will determine a number of consequences, i.e., propositions for which the truth value will differ from the current state. It is by (subjectively) evaluating these outcomes and by placing a personal preference on the *utility value* thus assigned that a player decides on the (subjectively) most favourable action. The consideration of personal preferences is where this approach differs from the one outlined in Simon's work (cf. [16]).

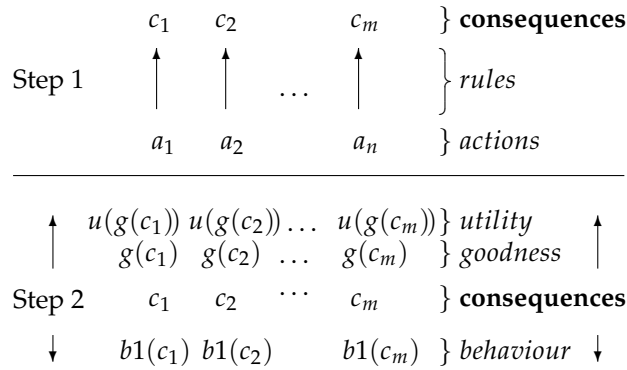
We define a strategy  $\sigma$  formally:

**Definition 14** (Strategies  $\sigma$ ). Let strategy  $\sigma$  be determined by  $\langle g, u, b1, b2, \succ \rangle$ :

- $g$  a function  $g : \mathcal{C} \rightarrow (\mathbb{R} \times \dots \times \mathbb{R})_n$  mapping consequences to multi-valued goodness values representing specific aspects of these consequences (such as reaching the specific goal of, e.g., taking a pawn).
- $u$  a function  $u : (\mathbb{R} \times \dots \times \mathbb{R})_n \rightarrow (\mathbb{R} \times \dots \times \mathbb{R})_m$  mapping goodness values to utility values (representing, for example, how taking that pawn will serve one of a number of strategic objectives, which ultimately lead to winning the game). The arity of  $u$  and that of its output may differ.
- $b1$  an evaluation function  $b1 : \mathcal{S} \rightarrow (\{0, 1\} \times \dots \times \{0, 1\})_k$ , which maps a state to  $k$  boolean values. Each of these values indicates whether a formula  $\phi$  is true in that state (i.e., whether the formula is valid, given the valuation (assignment of truth values to propositions) for that state). Each of these  $k$  formulae represents a behaviour which we want to support.
- $b2$  a function  $b2 : ((\{0, 1\} \times \dots \times \{0, 1\})_k \times (\mathbb{R} \times \dots \times \mathbb{R})_m) \rightarrow (\mathbb{R} \times \dots \times \mathbb{R})_m$ , that combines the  $k$  boolean values calculated by  $b1$  and the  $m$  utility values. The output are  $m$  values which combine the utility as well as the behavioural preference of the action.
- $\succ$  a preference relation over  $\mathcal{C}$  such that  $c_1 \succ c_2$  iff  $b2(b1(c_1), u(g(c_1))) > b2(b1(c_2), u(g(c_2)))$ .

We then define  $\sigma : \sigma(\mathcal{A}') = a_i$  iff  $\forall a_j \in \mathcal{A}' \setminus \{a_i\} : r(a_i) \succ r(a_j)$ . Finally we introduce  $\Sigma$  as the notation for a set of strategies  $\{\sigma_1, \dots, \sigma_n\}$ .

Figure 3 illustrates this: From a set of available actions  $a_1, \dots, a_n$ , the player can bring about a set of consequences  $c_1, \dots, c_m$ . The player can then assign goodness values to each of these consequences and subsequently evaluate them according to his utility estimates. This allows him to order the consequences according to his preferences and thus to pick the most favorable action.



**Figure 3.** The proposed model for automated rational behavioural decision making by an artificial intelligence. The preference over some actions is calculated by considering the consequences  $\mathcal{C}$  such that  $c_1 \succ c_2$  iff  $b2(b1(c_1), u(g(c_1))) > b2(b1(c_2), u(g(c_2)))$ .

## 6. Proof of Concept Implementations

### 6.1. Aims and Objectives

Two main claims are made regarding: (1) the proposed formalism for behaviour; and (2) the modelling for behaviour driven rational AI on the basis of this formalism. These claims are:

1. The formalism is suitable for controlled environments such as simulations and computer games (assuming that their data structures are designed appropriately). The evaluation and comparison of formally stated behaviours, as well as the translation thereof into their natural language equivalent, is straightforward and can be automated. The algorithms for doing so are computationally efficient and scale well.

2. Using formal behaviour statements (expressed in our formalism), we can augment standard models for rational decision making from the literature to include *behavioural stances*. Using this model, we can design and implement a game-playing AI whose choices exhibit clear (and human-like) behavioural preferences.

To validate both claims, two separate games were designed and implemented/realized to evaluate our approach and to *play test* the resulting serious games. The reader should keep in mind that the objective for both of the implemented games was to validate our work and to drive improvements by allowing insights. Neither of these games is a full game that can be released to the public in its current state (but both could be: the work required to finish either, while substantial, is straightforward and no unsolved issues prevent a completion of the implementation).

Regarding the design of AI players with human-like behavioural preferences (promised in claim 2): we evaluated the behaviour exhibited by human players in the card-board version of the game *SoxWars* (described on p. 36). Admittedly, due to the time required to test-play this version of the game, only very simple behaviours were identified. However, these suffice to show that the AI can be designed to adhere to formally defined behavioural stances. The behaviour implemented as a proof-of-concept for the mobile phone version of the game (described on p. 39) is “*play competitively against anyone who has played competitively (against anyone) in the last round*”.

## 6.2. Proof of Concept Game—Utility Tycoon

### 6.2.1. Objectives

To validate claim 1 of Section 6.1, a game was designed around two opposing types of behaviours, driven by a stance either for or against *green energy*. Please note that this is explicitly a proof of concept implementation and *not* in any way a representative statement about what constitutes green energy. Integrated in a resource-management game, the players could opt for either behaviour through their choices within the game (cf. Figure 4 for screenshots), with the assumed less preferable option (coal and nuclear power plants) being slightly more advantageous in the game. At an advanced stage in the game, the supported option (renewable energy sources) provided small incentives and *bonus events* such as good PR and romantic interests (cf. Figure 5 for screenshots).

The objective was to validate the statement that a game can be designed to realize atomic actions which can be clearly identified as belonging to one of the two subjective stances. In addition, complex behaviours within the game can be formally expressed. An automated tool was implemented to show that the translation of complex formal behaviour statements into natural language and back is possible. Through a *moderator tool*, behavioural statements (which control the unlocking of the bonus events shown in Figure 5) can be changed during the course of the game (cf. Figure 6 for screenshots).

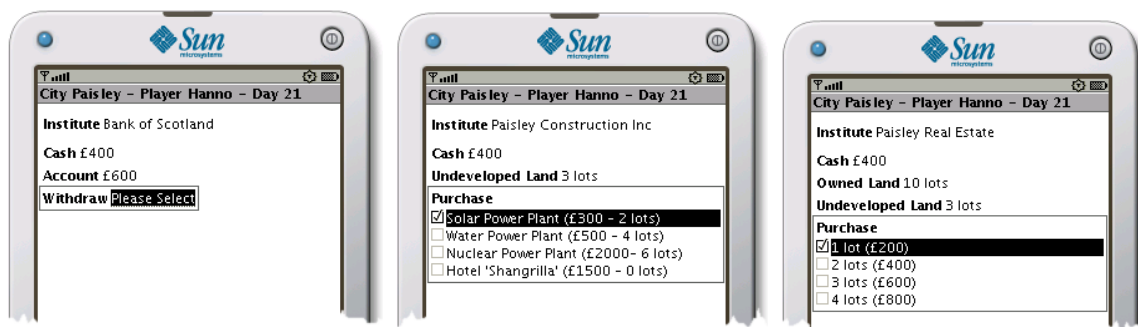
To validate the computational efficiency of the approach, as well as to showcase that such games do not require excessive resources, the game was implemented for a mobile phone (using an emulator tool, see Section 6.2.3 for details on the implementation). Minimal performance specs were assumed and the available space to display the game was considered to be minimal.

### 6.2.2. Brief Description of the Game

The game is a typical resource-management game (cf. Section 4.2). The player is starting with a restricted amount of money and is given a number of options to start a company, situated in the utility sector. The game can include interaction with the game AI, for example by engaging in politics in the game. In addition, as shown by the second game, the game AI would be able to follow behavioural guidelines.

In *Utility Tycoon* [145], the player assumes the role of the CEO of a company/start-up that produces and sells utilities in a fictional country. The player is competing for market shares in different cities, and the products sold are water, electricity and gas. The resources include the land to build the production sites on and the infrastructure to transport the utilities. Ideally, the player manages to corner the market in at least one city, if not in all (see Figures 4–6 for screenshots).





**Figure 4.** Game screens on an emulated mobile phone display: handling resources such as money (**left**) and real estate (**right**) as well as making behavioural decisions by deciding for or against investing in sustainable energy production (**middle**).

### 6.2.3. Game Design and Implementation

Neither the design of good games nor the implementation of computer games in general is the topic of this article and no background is provided here. We briefly discuss the technical details of the implementation as well as the design decisions that were considered for the game.

#### Serious Game Principles

In [146], we discuss 10 principles of good serious games in relation to RMGs. Our game meets these as follows: the player takes on the role of the person in charge (**Identity**). The production and use of the resources is directly dependant on the decisions and actions of the player (**Interaction**). Since the core of the game is independent of its appearance to the user, it can be quickly adopted to appeal to players in a large variety of scenarios. In previous work [145,147,148], we have proposed a customizable RMG built on a set of formally stated behaviours. In earlier work [149], we reported on a framework that facilitates personalization of a game to tailor it to the needs of the individual player (**Challenge and Consolidation** and **Customisation**). The course of the game is directly influenced by the decisions of the player and immediate feedback is provided by the game (**Production**). The game provides a familiar setting and the player is encouraged to make decisions that could have dramatic results in the real world. While the effect of taken actions can be very drastic within the game, no real world consequences whatsoever exist (**Risk Taking**). An individual game can take a long time but consists of many re-occurring themes and tasks that are repeatedly faced by the player. The game designer can design the game in such a way that the player can always do well but has to excel in order to rise to the top, either through a well tuned AI or though the game dynamics (**Pleasantly frustrating**). The problem and all relevant aspects thereof are made available to the player in an understandable description, and the player can improve by making use of the provided information (**Well-Order-Problems**). The game is constructed such that problems of the same type can be solved in an analogous manner throughout the game. This is an incentive for the player to abstract (**System thinking**). The player is directly responsible for failures and success (**Agency**).

#### Implementation

The tested version of the game is a prototype that was implemented using the Eclipse SDK (Version: 3.3.1) and the Java Wireless Toolkit (JWT) 2.5.2 for CLDC with the device configurations set to Connected Limited Device Configuration (CLDC) 1.1 and Mobile Information Device Profile (MIDP) 2.0. For the performance evaluation, we created statements that were artificially designed to require the longest possible evaluation (e.g., we applied the opposite of the suggested normalisation, cf. Section 5.3.2), i.e., we tested for the worst case scenarios. The performance of the prototype implementation was tested in the Java Wireless Toolkit (JWT) 2.5.2 for CLDC using the above settings.



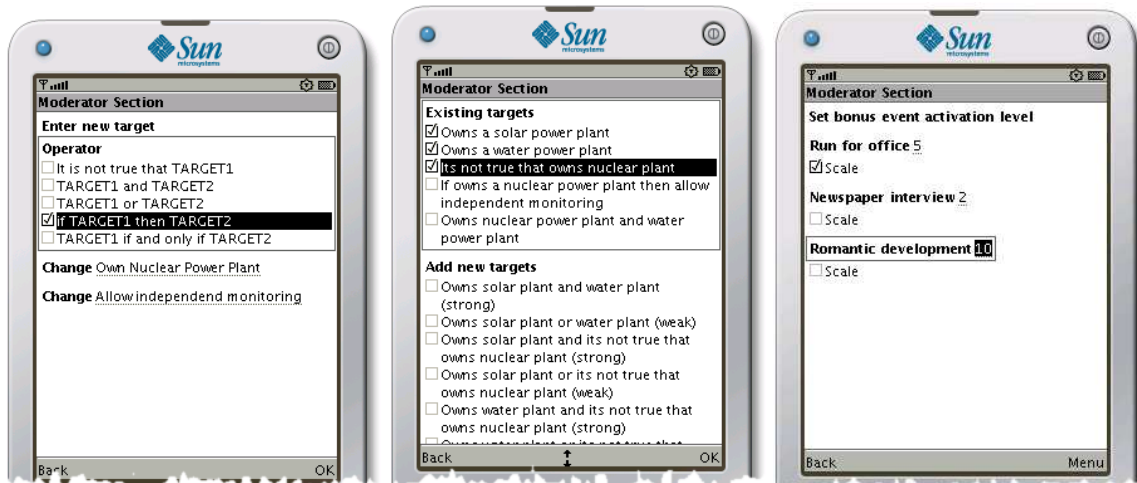
**Figure 5.** In line with wanting the game to be *fun*, bonus events are included. Participating in these is not considered a regular action of the player and thus is not evaluated as behaviour; however, the player can unlock these events through adhering to specific behaviours.

## 6.2.4. Validation of Our Approach

### Representation and Evaluation of Formalised Behaviour Statements

Underlying the game are formally stated behaviour statements [147]. The state of the game is evaluated automatically [148] by an evaluation module. At any given time, the game can report on which of the defined behaviours the player has and has not exhibited. In addition, any newly created behaviour description can be evaluated against the game as far as it has been played so far.

The moderating tools (see Figure 6) are an important feature of the game as they allow the adoption of the behavioural statements as well as the thresholds for the bonus events during gameplay. Due to the manner in which this is designed, the user does not have to have any prior experience with programming at all. Not the least because of the restricted screen size of the average mobile phones, all interfaces are kept simple and their functionality is straightforward.



**Figure 6.** Emulated mobile phone displays showing various interfaces for the game designer: (left) a simple interface allows the construction of complex statements from simpler ones, offering the five connectors defined in Section 5.3; (middle) the activation of newly constructed statements; and (right) the setting of a threshold of how many such behaviours have to be exhibited to unlock a specific event.

### Computational Efficiency and Performance

For the evaluation of behavioural statements (for which complexity increases linear with the length of the target), we tested for 50 statements consisting of 10 or 20 atomic statements each, and report the average over 50 tests, each comprised of 100 runs. The simulated phone averaged 0.019 ms (length 10) and 0.025 ms (length 20). When testing the statements for satisfiability, we used the

same statements as before but averaged over 50 test comprised of 10 runs each. The prototype was able to check satisfiability for all behavioural statements in 431.18 ms (length 10) and 439.05 ms (length 20). The very small increase in computation time is due to rounding and the fact that the execution was often too fast to register at all, i.e., the execution took less than 1 ms. These results show that, even for unrealistic length of behaviour statements (consider that sentences containing 20 individual facts are rarely used in real life) and for large number of statements, the proposed formalism can be used: checking for satisfiability will take less than half a second. We argue that this shows that the approach is of sufficiently low computational complexity to be of general use.

### 6.3. Proof of Concept Game—SoxWars

#### 6.3.1. Objectives

To validate claim 2 of Section 6.1, a multi-player resource-management game (where all but one player are played by the computer) was designed as a *pen-and-paper based* board game. The focus of the game is on cooperative and competitive behaviour, specifically in the context of the behaviour previously observed from the other players. The game was designed with the aim to create specific recurring situations to which the players can relate and where the context (the opponents' previously exhibited behavioural stance) is accessible to the player. We motivate our choice for competitive versus cooperative behaviour by the fact that this type of behaviour has seen enormous interest from the field of behavioural psychology. The *Prisoner's Dilemma* [14,41,51] described in Section 2.3 revolves around this type of behaviour. We argue that in the context of designing formalisms and models for rational behavioural AI, these relevant blocks of interaction will serve an important role when building artificial cognitive systems capable of interacting with humans. The ability to adapt the AI behaviour to behaviours observed from the individual human is crucial in this context.

#### 6.3.2. Brief Description of the Game

In *SoxWars*, a player competes directly for resources with the other players in order to first conquer, and later control, segments of a finite market capacity. As the game adheres to serious game principles in a very similar way to *Utility Tycoon* (cf. Section 6.2.3), we do not repeat this paragraph here. Suffice it to be said that the game follows the basic elements of resource-management games (cf. Section 4.2). In short:

- The players initially start with a small amount of money (resource).
- Using that money they can purchase supplies of socks for stock (products).
- The shops where these socks can be sold are limited and there is a system in place that favours the supplier that has, in the past, supplied the respective shop.
- The game is turn-based, and turns consist of a number of phases, the order of which is fixed.
- The revenue from selling socks is fixed while the cost for restocking (i.e., the acquisition of socks) varies depending on the phase of the turn when it happens.

Furthermore, in the context of using the game to evaluate player behaviour, the following holds:

- During each phase, players make their choices simultaneously. These decisions, which can affect the outcome for all players, are being revealed directly afterwards (i.e., before the next phase).
- The game is designed to converge to situations where trade-offs are required. While a balanced and mutually fair distribution of opportunities is possible, any player can upset this balance and force the game into a series of conflicts (i.e., situations where players will compete for something).

Two limiting factors contribute to the dynamic of the game: firstly, the total number of products that can be sold per round is limited to the total number of streets of the town. Secondly, the ability to sell a product in a street is directly related to having supplied the shops in that street in the past. Due to this, it becomes strategically important to sell in a certain number of streets, which in turn means that

the tactical decision to purchase the product is not only dependent on the price for the product but also on the number of products that are required to achieve and sustain this strategic objective.

Furthermore, the total amount of products available for purchase (i.e., to resupply the player's stock) in each round is also limited. They are offered to players in equal parts but there is one phase during each round where players may bid on the supply offered to another player. This is made less attractive by the fact that bidding on another player's products happens for a price that is actually above the fixed sales price, i.e., inevitably incurs a financial loss. As it is crucial for the effort to progress in the game to have more than the allocated average amount of products, bidding on another player's products becomes a strategic decision which, while necessary to improve one's ranking, is not one taken lightly due to the financial loss. The decision on whose products to bid is equated to the behavioural decision of whom to *support* and whom to *attack*.

### 6.3.3. Game Design

The game progresses for all players in the same way. The available decisions are identical and are taken at the same time. Due to this, if all players simply sell the products that they can buy without interfering with the other players, the game will settle in a state where everyone is reaching the same conditions and the streets are effectively evenly distributed between the players. Using a neutral/cooperative playing AI, the game can be used to evaluate the human player's behaviour as either competitively or cooperatively. This evaluation can be performed in an automated manner on the basis of a set of observed and formally stated behavioural statements (cf. Section 6.2.4 for *Utility Tycoon*).

In addition, such formalized behavioural statements can be used in the design of the AI player. An AI can be instructed to adapt their behaviour in general (for all players), for specific opponents (e.g., *play competitively against Player B*) or even in reaction to observed behaviour (e.g., *play competitively against anyone who plays competitively against Player B*).

Because of this, the AI players can be used to control the dynamics of the game. In addition, a carefully designed set of AI players can basically ensure that certain situations will occur during a game. This means that the game can be used to assess human behaviour in very specific situations.

#### The Different Phases of a Turn

The game is a turn/round based game, i.e., it repeats rounds during which the players take actions. Each round has three types of actions: resource-acquisition (RAC), resource-assignment (RAS) and resource-allocation (RAL). There are three phases of resource-acquisition and -assignment before the round ends with the allocation of the players' resources:

**RAC** A limited number of new products can be acquired per round in three separate RAC phases:

- RAC1 Buying: Each player is offered the same number of socks, for \$1 per unit.
- RAC2 Bidding: The players are offered additional socks at a cost of \$1.50 per product. Players can also choose to *bid* on the products offered to other players (at the inflated price of \$2.50).
- RAC3 Trading: Players can offer remaining resources to other players for \$2 (market value).

**RAS** There are a number of territories, each with a number of shops where the resources are sold for \$2 during RAL. Assignment happens in three phases and only to territories, not to specific shops:

- RAS1 Shops only accept resources from the player that delivered to them in the last round.
- RAS2 Shops that had a supplier last round only accept resources from that player.
- RAS3 Delivery to any remaining (not yet supplied) shop.

Conflicting deliveries are handled during the RAL part of a turn.

**RAL** There is a bias towards players who supplied shops in the last round, making it beneficial to reliably supply *your* shops. This is especially relevant since the number of shops is finite.

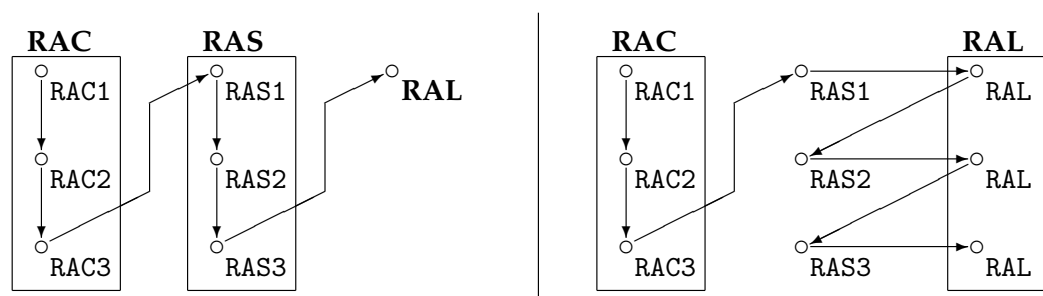
Shops first accept delivery from players that supplied them in the last round but then accept supplies evenly from all supplying players (on a territory by territory basis). Conflicts are resolved by random allocation such that all players are favoured in turn. If a fair (but random) distribution is not mathematically possible, the human player is favoured (by design).

Essentially, it is in the players' interest to deliver (and keep delivering) to as many shops as possible. This will let the game converge to a state where all shops are loyal. At this stage, the game will produce only the number of resources required to satisfy the demand of all shops. Once this happens, the only way to acquire more territory is to bid on another player's resources (at a loss) in the hope of keeping the shops this player can no longer deliver to in the next round.

### Three Variations for the Order of Phases in a Turn

Three different models for the order in which certain phases would occur in a turn were considered and, to a certain extent, tested. Specifically, in a very basic form, the game was test played with the three different ordering of the phases. The motivation for these models is straightforward: the game was designed to assess behaviour and the individual phases were created to contain certain interactions between the players and the game world. Some of these interactions have an affect on the other players, some only effect the state of the world. The implementation of an individual phase is a relatively modular task, meaning that phases are for a good part exchangeable. Therefore, the full game is basically calling the individual phases in a specific order, and that order can be changed and adapted without the need to adapt the phases themselves. Three variations were considered:

In the first model (Figure 7, left), the players acquire all resources (and offer their resources to each others in RAC3) before assigning them. While this makes for the nicest game-play, it stretches the feedback loop regarding the actions of other players. Depending on the information displayed, it might not be until the end of a turn that some of the consequences of an action become apparent.

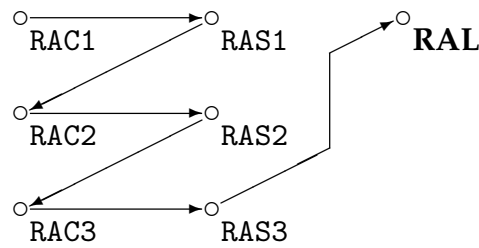


**Figure 7.** Two variations of a turn in SoxWars: (left) Model 1 where resource acquisition (all three phases) is followed by resource assignment (all three phases) before the resources are allocated; and (right) Model 2 where the resource acquisition is finished before any assignment happens, but assignment and allocation alternate, allowing players to see the result of their assignments.

The second model (shown in Figure 7, right) alternates the resource-assignment and -allocation phases. This emphasises the different choices in the assignment phases, as conflicting situations may arise (in which case some players retain their resources), drawing attention from the acquisition phases.

The third model (shown in Figure 8) focuses more on the resource-acquisition phases than on the resource-assignment and -allocation phases. Since the resource allocation is happening after all players have made their decisions, it becomes harder to draw conclusions about the action of other players.





**Figure 8.** Model 3: Resource acquisition and assignment are alternating before finally being allocated. This version puts the most pressure on the resource allocation phase as the impact of the supply as well as the tactics of the other players become more visible.

It should be noted that the models for the different phases (cf. Section 6.3.4) remained unchanged for the three different models of the game. Do to this, all three variations could be modelled very similarly with minor changes for the assessment of behaviours or the behavioural AI.

The game was implemented in three different formats (cf. Section 6.3.5): (1) on cardboard as a pen-and-paper implementation; (2) as a mobile phone based game; and (3) as a web-based game.

### 6.3.4. Modelling the Game

#### Modelling Resource-Acquisition

In the implementation of the game, the process of making resource acquisition decisions is serial, meaning that the game cycles through all players one by one. However, these decisions are not revealed to the other players until the end of the phase. In a game with more than one human player, this might be a design issue, but since all but one players are computer players, the issue of hiding information which is actually provided on the screen does not arise.

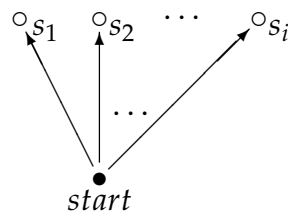
- **Modelling RAC1:** Phase one of the resource acquisition (RAC1) does not contain any behaviour of interest to us. Whether a player decides to purchase resources does not appear as part of the considered behavioural statements. Obviously, there are changes in the state of the game, but these can be represented as a single world in the model: if a specific player purchases resources, this only affects the propositions related to this player's stock and funds. These propositions are disjoint with the propositions of the other players and thus we can express all the changes in  $\Phi$  within a single world. The only minor liberty taken in this approach is the fact that the globally available resources are of course decreasing every time a player purchases stock, which happens multiple times in the stage. However, global resources are not considered for our behavioural statements as they are not under the control of the player. Therefore, they are not included in  $\Phi$ .

The difference in the frames (i.e., the models without the propositions assigned) for the stages is thus mainly in the number of possible future states (e.g.,  $i$  in Figure 9:  $s_1$  to  $s_i$ .) This means that if the player can decide on the number of resources to purchase, and if there are exactly  $n$  products offered to each of the  $j$  players, there are  $n + 1$  different results for each player (allowing for zero products being purchased), resulting in  $i = j \times (n + 1)$  in the model for RAC1. In our implementation, we included this information as it was relevant for the expression of the rational aspect of the AI; however, we restricted this decision to "buying" and " $\neg$  buying", so that in our implementation  $i = j * 2$ .

- **Modelling RAC2:** The second stage of the resource acquisition is the most important one for the behaviour analysis. Again, the model can be collapsed to the model shown in Figure 9. This time, however, we consider the actions of the individual players with regard to the other players as the bidding on resources happens by one player but targets the resources of a specific other player. As above for RAC1, we only allowed the bidding on resources and did not enable a quantification for this (i.e., it is not possible to bid on a few resources of a player, it is either bidding on all offered

resources or bidding on none). We furthermore did not offer the option to “*bid on all players*”, forcing the player to select every opponent individually. We furthermore required that one bids on one’s own resources before bidding on those of other players. This is rational strategic behaviour and removes a number of complex behavioural constructions such as bidding on other players’ resources at the cost of not bidding on your own (which would be cheaper). This means that for  $j$  players there are  $j - 1$  other players to bid on, “*bidding only on the resources offered to the player*” and “*not bidding at all*”. Due to this, there are  $j \times (j + 1)$  possible combinations, and thus in the model for RAC2  $i = j \times (j + 1)$ .

- **Modelling RAC3:** In the last stage where resources can be acquired, we ignored the decision to accept resources offered. The justification for this was that including this increases the complexity of the represented behaviour by allowing for *sulking* and other emotional responses. The main justification for being able to omit these more complex behaviours is that the AI players will make the decision to purchase such resources on a purely tactical basis. The idea is that the human player is aware that the opponents are played by a computer and it is assumed that emotional responses are not exhibited towards these players. The decisions in RAC3 are very similar to the one in RAC2, in that each player gets to decide whether to offer a fixed amount of stock to another player. Therefore,  $i = j \times (j + 1)$  for model RAC3 as well.



**Figure 9.** If there are no interactions between the individual choices made, serial decisions of players (which result in a model with a number of layers) can be collapsed into a model with a root and a number of states (which can be reached in one step from the root) without loss of expressivity.

### Modelling Resource-Allocation

The three allocation phases are not as well modelled as the acquisition phases. Early into the game design, it became evident that, while there was a great potential for conflict and interesting game dynamics in this aspect of the game, and it had very little to do with the behaviour under investigation. Therefore, the same approach was taken as for the RAC stages: the decisions of the players are hidden from their opponents until after the stage is completed (see Figure 9 for the resulting mode).

The RAL stages are used to remove as many conflicting allocations as possible with the aim to remove as many conflict resolutions from the final stage, RAL. This is achieved by first ensuring the save allocation to shops in RAL1. During this stage and the following stages, only the shops that are available for delivery are highlighted (this was as much an interface design choice as it was a simplification designed into the stages intentionally). This is followed by RAL2, where the opponent’s choices from RAL1 are revealed, but only in the form of shops not being available for delivery, without an indication of which player has allocated their products to them (though in most cases this is obvious). Finally, in RAL3, it should be obvious where an allocation has no chance of success, which should make the remaining allocations less likely to result in returned stock. The aim was to contain the player interaction to the RAC2 (and to a lesser extent to the RAC3) stage as much as possible.

### Modelling Resource-Assignment

The AI players do not consider the possibility of a conflict resolution. The argument is that this does not affect their game-play and, since these are random decisions, does therefore not constitute a behavioural decision. As the allocation of resources to a shop is not considered competitive behaviour by us, this simplification does not restrict the AI’s ability to be guided by behavioural stances.

### 6.3.5. Implementation

#### Implementation: Cardboard Version

The game was first designed on paper and with a clearly defined set of behavioural tests in mind. It was then implemented and tested as the board game (between human players only, some of which were sometimes instructed to play in a certain way) shown in Figure 10. This was used to provide early evaluation of the design and game mechanisms. This led to a number of minor revisions.

The motivation for implementing the game as a cardboard based game was that this was inexpensive and allowed for rapid adaption if changes needed to be made. As mentioned above, the game was developed on the basis of a set of target behaviours. As such, it was heavily revised in the early stages until it reflected the intended behaviours well while at the same time containing as few unnecessary elements and aspects as possible to avoid distracting from the behaviours.

The aim was to decide on the individual stage and to investigate the feasibility of constructing a resource-management game (cf. Section 4.2) around the intended tool. Once the overall idea of the game was decided upon, the separate stages were designed and, as discussed above, their order considered.

Furthermore, through the initial test play, a number of variables and parameters were tuned to bring about a specific dynamic in the game. These parameters were only investigated superficially, that is, to the point where the game was flowing and the rough direction was guaranteed. This was not further fine tuned throughout the remaining development process, as the game was never intended to become a fully playable game on a par with commercial games.



**Figure 10.** The laminated paper version of the game. White board markers can be used to update the individual scorecards as well as the playing board. On the playing board, the seven streets are represented by hexagons, each of which has six shops represented by isosceles trapezia (wedge shaped) which can be covered by cut out paper tokens to represent the player who currently controls it.

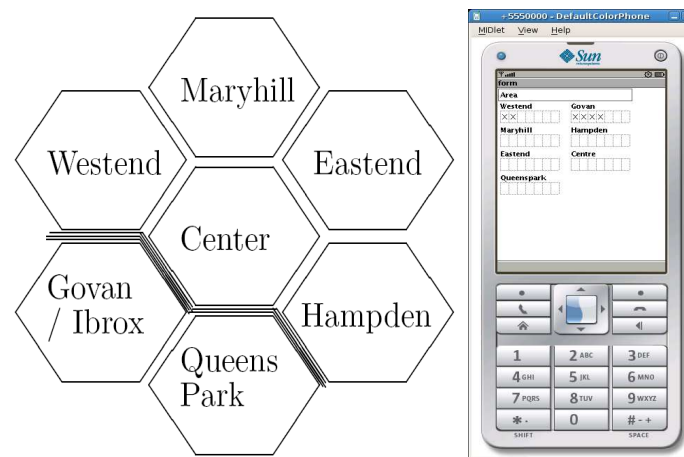
#### Implementation: Mobile Phone App

Following similar considerations as for *Utility Tycoon* (cf. Section 6.2), the computer game version of the game was developed for emulated mobile phones (see Figure 11). The reasoning behind this choice is that: (a) by implementing the game for low specs devices and with very limited graphics options and screen space, we can motivate the claim that the approach is computationally efficient and scalable; and (b) the intended use of the game as an evaluation tool requires that it can be deployed on platforms that are both widely used as well as available in the contexts when humans are willing to engage in casual play. This makes the implementation for mobile devices an obvious choice.

The formalized statements were written in Prolog format, mainly for historical reasons. There had been the initial aim to implement the evaluation algorithms using Prolog, but that approach was soon abolished as the required functionality could directly and very easily be implemented in using Java.

The technical details are the same as for *Utility Tycoon*, discussed in Section 6.2.3. The test-playing of the game was also conducted exclusively using the emulator. In other words, the game was never tailored for a specific phone type and played on a real device. The aim for this series of tests was to identify bugs and to test the ability to generate records of the human player's behaviour. In addition, the AI was implemented and active during the test games, for which colleagues were recruited as testers. At this stage, the game was neither visually attractive nor particularly *fun* to play.

While the paper implementation allowed the simulation of the AI through a human and on paper, a computer based implementation was needed to validate the approach. The intentionally restricted behavioural scope of the actions in the game helped to design the required formalisms and data structures. This was further facilitated by the low complexity of the game which allowed for a relatively direct implementation of greedy strategies for the rational AI and, regarding the behavioural AI, for a simple evaluation of histories paired with the consideration of the few possible actions.



**Figure 11.** The implementation for emulated mobile phones: the board with the individual parts of the city (left); and a screenshot (from the emulator) of the resource acquisition phase (right).

#### Implementation: Web-Based Game

Less for the evaluation of the approach but in line with the aim to deploy the game as an assessment and evaluation tool for large research experiments, an implementation for web-based platforms like a social web site was done (Figure 12). This is to show that the formalism can be implemented in applications that run efficiently on the platforms that have emerged as one of the current trends for games. Games such as *Farmville* attract millions of players, suggesting that a visually interesting but otherwise simple game can potentially capture the attention of a large number of users.

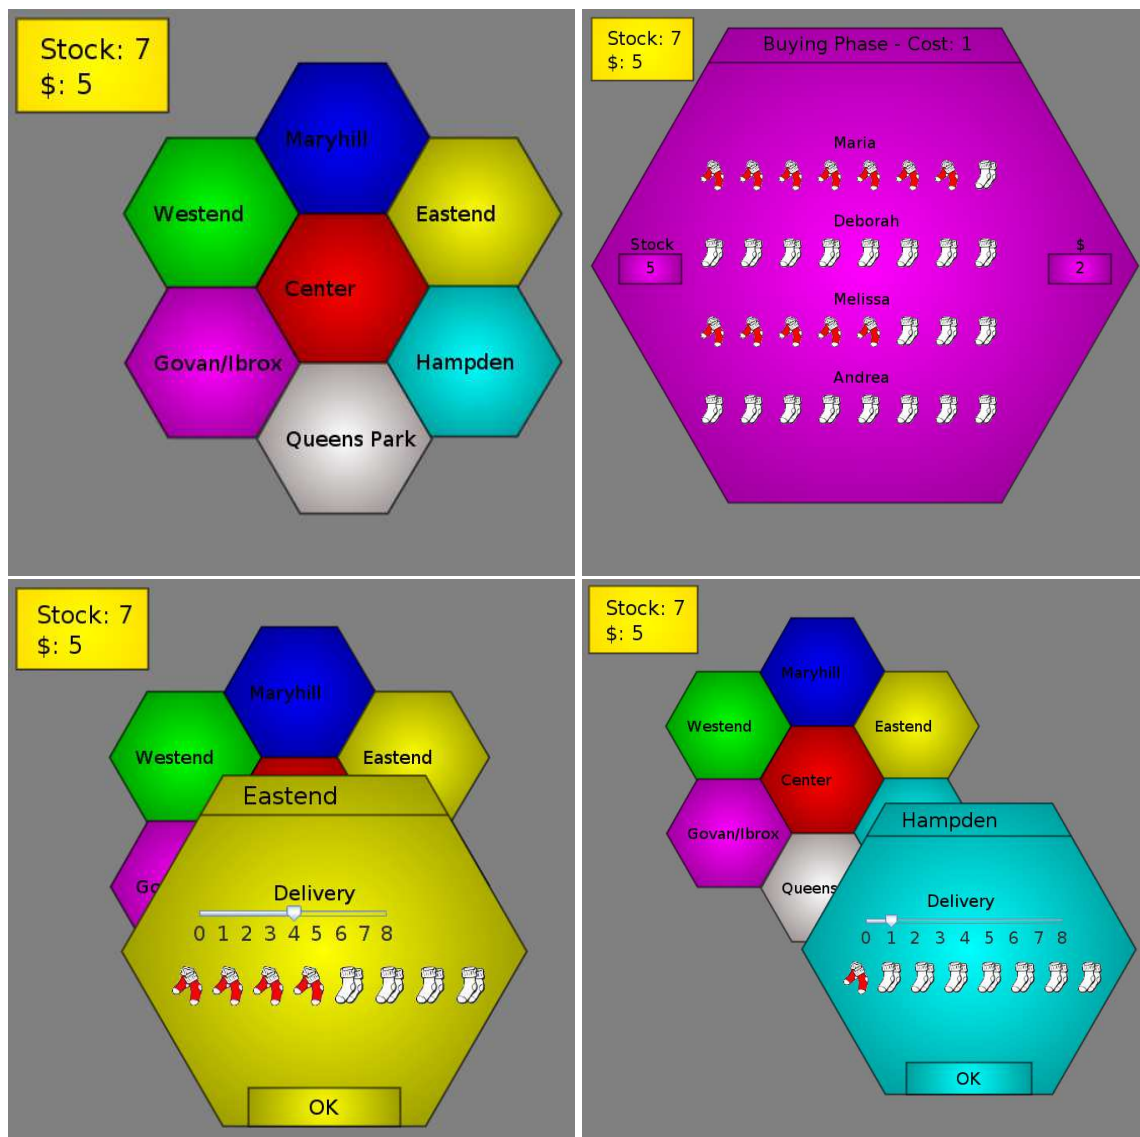
#### 6.3.6. Validation and Evaluation of Our Approach

The validation of the approach is not entirely disjoint from the evaluation of the game as an attractive means to pass the time (i.e., to engage in play). The three implementations were tested to various levels but only the mobile phone based game was tested for the performance of the AI.

#### Lessons Learned: Cardboard Version

Regarding the game-play: The game was play-tested by four colleagues only (see Figure 13). This was sufficient to have other humans provide initial feedback on the framing story of the game, on the overall idea and to verify that the game could be explained in a short and precise way.





**Figure 12.** Screenshots for the user interface for the web-based version of SoxWars, developed for interactions using mouse, mobile phone keys as well as touch screen input: the board (**top, left**); the screen for buying products (**top, right**); and the delivery decision to two territories (**bottom**).



**Figure 13.** Play-testing the paper version of the game (cf. Figure 10). One player calculated all his actions using the model for the game as well as our algorithm for rational behavioural AI. While slow, this allowed the observation of AI game-playing behaviour and the analysis of the exhibited behaviour.



As there was neither a mechanism in place to automatically record behaviour, nor the means to use formally described behaviour to steer the players, the colleagues were instructed to play with a certain behavioural stance. This led to a number of observations and comments regarding aspects of the game as well as the order of the stages. For example, the fact that conflict resolution was performed by rolling a dice meant that maintaining a certain behavioural stance became a matter of interpretation of whether another player had intended or even foreseen the outcome of the random event. While this is certainly something that can be considered, it was not what was intended for the game at hand. This and other insights drove the decision for the overall model for a turn (as described above in Section 6.3.3).

#### Lessons Learned: Mobile Phone App

The game was tested and evaluated by a dozen undergraduate and postgraduate students and research staff. It should be noted that at this stage the game was a single player game, that is, only one human was required to play, as all other players were controlled by the computer. During the initial testing, which amounted to not much more than random choices by the human tester to go through the game, many bugs and issues with the AI were identified and solved. This led to the first playable game. As the testing progressed, we first finished the implementation for the rational AI, then finalised the logging and recording of histories and finally used these records to enable the AI players to influence choices by their behavioural predispositions. The rational AI was relatively straightforward as the evaluation of the choices according to how well they fit in with the game is almost trivial. This is of course by design, as there is very little difference in claiming a shop in one street over claiming another shop in another street. The rational player aiming to win will attempt to acquire as many extra resources as possible, with the only side consideration being on the ability to subsequently *defend* the newly claimed territory (which is a matter of having enough money to secure the required resources). However, once the formalism to express behaviours was in place, we added the behavioural stance to the rational decision making and thereby created different AI players.

With regard to the evaluation of our formalism, the aim was: (a) to implement behavioural stances; and (b) to verify through play that these stances were honoured by the AIs. This had practical implications on the complexity of the implemented behavioural stances: we implemented, e.g., *“play competitively against anyone who has played competitively (against anyone) in the last round”*. We did not implement *“... in the last five rounds”* or *“against anyone who played cooperatively”*. The argument is that the tested example requires considering other players’ behaviour (which is what we had set out to do) while avoiding complexity issues. In addition, to verify the AI behaviour we had to play multiple rounds, by considering behaviours with deeper nesting of statements would have required playing more rounds, in considerably more games. The tested behaviours suffice to show that the formalism works and that it can be used to drive behavioural AI.

In contrast, the evaluation of the implementation in the context of it being a playable game focused on the functionality of the prototype. Other than in testing the automated verification methods and the ability of the game to bring about certain situations in the game, the recorded histories were not evaluated with respect to the behaviours exhibited. By that we mean that we did not interpret the collected data to evaluate the players from a psychological point of view. The game was never tuned by an expert with a background in psychology and the specific choices made in the formalisation of the behaviour (described above) would likely raise objections by the practitioners. The prototype was implemented to test the methodology and to showcase the formalism.

#### Lessons Learned: Web-Based Game

The web-based version of the game was only implemented rudimentarily because the formalism on which it would be built is identical to the mobile-phone based version (i.e., the same data structures and methods). Therefore, only the interface was developed and based on simulated data structures

to show that the user interface from the mobile phone version could be improved for a web-based version. This implementation was not playable.

As the design of a game requires a lot of time and dedication, and due to the view that in order for the game to become an appealing and *fun to play* game, a lot more time had to be invested in a variety of matters that have little to do with the proposed formalism and the aim to design a behavioural AI, no further effort was extended towards the web-based game. We report on it here to suggest that a properly designed game might be best implemented for a web-based interface as the ultimate requirement will always be to attract players to actually engage in play.

## 7. Summary and Conclusions

This article proposes to use computer games for the assessment and subsequent evaluation of human behaviour. This—in our opinion—requires a formalism to express behaviour as well as models for how behaviour comes about. Given these, we argue that computer games are ideally suited to provide comparable and controllable test environments because they are: (a) discrete (at any time their exact state can be recorded); and (b) entirely controllable. With regard to the latter, many games require a number of players (and human behaviour is often evaluated in the context of/in reaction to the behaviour of others) and therefore we propose an approach to design AI players that are bound by behaviour directives expressed in the same formalism. The presented work is based on established models from the field of behavioural psychology, formal logic as well as approaches from game theory and related fields. The suggested model for behavioural AI has been used to implement and test a game, as well as AI players that exhibit specific behavioural preferences.

We discuss models for behaviour from the field of behavioural psychology but kept our work generic enough to argue that the model on which our implementation is based could be exchanged if one desired so (or if the specific application in mind suggested it). We introduce propositional logic as a formal language for the unambiguous expression of statements, and extend this by definitions for modal logic which allows concepts such as time (sequences of events) and different possible futures (multiple possible successor states). We show that—within a restricted domain—one can use logic formulae to express classes of behavioural choices/actions within a game. We furthermore discuss well-known approaches to model rational behaviour or social choices/actions. We suggest that we could extend a model for rational decision making by adding preferences over the outcomes of actions with regard to formally stated behaviours (or entire classes of so-stated behaviours). To motivate this claim, we designed and implemented a resource-management game and play-tested it against AI players. Our AI played rationally while adhering to the predefined behavioural stances and, when using multiple AI players, we could force certain game situations to occur.

This article provides the reader with an approach that can be tailored to suit specific needs. While this concludes the article, it is merely a beginning. The discussed games were designed by a computer scientist and meant as proof-of-concept implementations; future work should ideally be done in the context of well-designed experimental setup and commissioned/supervised by practitioners in the field of psychology. Several aspects, such as ethical approval, a well-formulated research question, etc., need to be addressed but we are confident that our work has the potential to contribute greatly to the field of behavioural psychology. Computer games are pervasive and, once a game is designed and implemented, the operational cost of deploying it to a large audience can be very low. As such, we see the potential for collecting large amounts of data across demographic boundaries.

Furthermore, the approach lends itself for use by the computer games industry. In this article, we use competitive versus cooperative behaviour as our running example and we believe that the ability to impose such different behavioural stances on AI players in commercial games can significantly add to the playing experience. The level of detail as well as realism is entirely up to the designer of the formalism; the possibilities for creating individual AI characters are virtually limitless. Moving away from scripted behaviour, our approach enables the design of quirky but consistent AI players that will act consistently even in entirely new situations. In addition, this could be used to implement

diplomatic elements of a game and could even lead to a new sub-genre, where the player has to figure out the AI's behavioural make-up to interact with it to achieve a specific goal.

The proposed approach is as much a guide line as it is a different angle on AI player design. It was never meant to be a complete solution but a demonstration of what can be done. Where we go from here is up to you, and the community. We respect the fact that we are not, in fact, experts in either psychology or computer games design. However, we would be thrilled to be contacted by such experts who are interested in taking our approach and using it in their fields.

**Conflicts of Interest:** The authors declare no conflict of interest.

## Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence	
DOAJ	Directory of open access journals	
ToPB	Theory of Planned Behaviour	Section 2.4
TACT	Target, Action, Context and Time	Section 3.1
PL	Propositional Logic	Section 3.2.1
ML	Modal Logic	Section 3.2.2
SEU	Subjective Expected Utility	Section 3.3.3
RMG	Resource-Management Games	Section 4.2

## References

1. Kitano, H. Intelligence in a changing world. *Nature* **2007**, *447*, 381–382. [[CrossRef](#)]
2. Turing, A.M. *Computing Machinery and Intelligence*; Oxford University Press: Oxford, UK, 1950.
3. Hildmann, H. *Computer Games and Artificial Intelligence*; Encyclopedia of Computer Graphics and Games; Springer International Publishing: Cham, Switzerland, 2018; Chapter 234.
4. Hassabis, D. Artificial Intelligence: Chess match of the century. *Nature* **2017**, *544*, 413–414. [[CrossRef](#)]
5. Silver, D.; Huang, A.; Maddison, C.J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **2016**, *529*, 484–489. [[CrossRef](#)] [[PubMed](#)]
6. Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. Mastering the game of Go without human knowledge. *Nature* **2017**, *550*, 354–359. [[CrossRef](#)] [[PubMed](#)]
7. Cho, A. 'Huge leap forward': Computer that mimics human brain beats professional at game of Go. *Science* **2016**. [[CrossRef](#)]
8. Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *ArXiv* **2017**, ArXiv:1712.01815.
9. Frenkel, K.A. Schooling the Jeopardy! Champ: Far From Elementary. *Science* **2011**, *331*, 999. [[CrossRef](#)] [[PubMed](#)]
10. Bowling, M.; Burch, N.; Johanson, M.; Tammelin, O. Heads-up limit hold'em poker is solved. *Science* **2015**, *347*, 145–149. [[CrossRef](#)] [[PubMed](#)]
11. Riley, T. Artificial intelligence goes deep to beat humans at poker. *Science* **2017**. [[CrossRef](#)]
12. Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; Bowling, M. DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science* **2017**, *356*, 508–513. [[CrossRef](#)] [[PubMed](#)]
13. Simon, H.A. What Computers Mean for Man and Society. *Science* **1977**, *195*, 1186–1191. [[CrossRef](#)] [[PubMed](#)]
14. Simon, H.A. The behavioral and social sciences. *Science* **1980**, *209*, 72–78. [[CrossRef](#)] [[PubMed](#)]
15. Behrens, T.E.J.; Hunt, L.T.; Rushworth, M.F.S. The Computation of Social Behavior. *Science* **2009**, *324*, 1160–1164. [[CrossRef](#)] [[PubMed](#)]
16. Simon, H.A. A Behavioral Model of Rational Choice. *Q. J. Econ.* **1955**, *69*, 99–118. [[CrossRef](#)]

17. Henrich, J.; Heine, S.J.; Norenzayan, A. The weirdest people in the world? *Behav. Brain Sci.* **2010**, *33*, 61–83. [[CrossRef](#)] [[PubMed](#)]
18. Hildmann, H.; Uhlemann, A.; Livingstone, D. Simple mobile phone based games to adjust the player's behaviour and social norms. *Int. J. Mob. Learn. Organ.* **2009**, *3*, 289–305. [[CrossRef](#)]
19. Hildmann, H.; Hildmann, J. A critical reflection on the potential of mobile device based tools to assist in the professional evaluation and assessment of observable aspects of learning or (game) playing. In Proceedings of the 3rd European Conference on Games Based Learning (ECGBL09), Graz, Austria, 12–13 October 2009; API, Academic Publishing International: Graz, Austria, 2009.
20. Hildmann, H. Mobile device based support to enrich initiative games and to facilitate in-game evaluation in experiential education. In *Proceedings of Techniques and Applications for Mobile Commerce (TAMoCo09)*; IOS Press: Amsterdam, The Netherlands, 2009.
21. Hildmann, H.; Boyle, L. Evaluating player's attitudes and behaviour through mobile device based serious games. In Proceedings of the IADIS International Conference on Mobile Learning 2009, Barcelona, Spain, 26–28 February 2009; IADIS Press: Florianopolis, Brazil, 2009.
22. Hildmann, H.; Hirsch, B. Raising awareness for environmental issues through mobile device based serious games. In Proceedings of the 4th Microsoft Academic Days, Berlin, Germany, 19 November 2008.
23. Hildmann, H.; Uhlemann, A.; Livingstone, D. A Mobile Phone Based Virtual Pet to Teach Social Norms and Behaviour to Children. In Proceedings of the DIGITEL '08: Proceedings of the 2008 Second IEEE International Conference on Digital Game and Intelligent Toy Enhanced Learning, Banff, BC, Canada, 17–19 November 2008; IEEE Computer Society: Washington, DC, USA, 2008; pp. 15–17.
24. Hildmann, H.; Guttmann, C. A formalism to recognise and monitor patient behaviour for intervention in intelligent collaborative care management. In Proceedings of the CARE Collaborative Agents—Research & Development Workshop at AAMAS 2011, Taipei, Taiwan, 2–6 May 2011.
25. Bandura, A. Self-efficacy: Toward a unifying theory of behavioural change. *Psychol. Rev.* **1977**, *84*, 191–215. [[CrossRef](#)] [[PubMed](#)]
26. Thorndike, E. *The Fundamentals of Learning*; Teachers College, Columbia University: New York, NY, USA, 1932.
27. Skinner, B. *The Behaviour of Organisms*; D. Appleton & Company: New York, NY, USA, 1938.
28. Skinner, B. *Science and Human Behavior*; Free Press Paperbacks; Macmillan: London, UK, 1953.
29. Bogost, I. *Persuasive Games: The Expressive Power of Videogames*; MIT Press: Cambridge, MA, USA, 2007.
30. Chomsky, N. A Review of B. F. Skinner's Verbal Behavior. *Language* **1959**, *35*, 26–58. [[CrossRef](#)]
31. Holden, C. Carl Rogers: Giving People Permission to Be Themselves. *Science* **1977**, *198*, 31–35. [[CrossRef](#)] [[PubMed](#)]
32. MacCorquodale, K. On Chomsky's review of Skinner's Verbal Behavior. *J. Exp. Anal. Behav.* **1970**, *13*, 83–99. [[CrossRef](#)]
33. Chomsky, N. *Language and Mind*; Harcourt, Brace & World: San Diego, CA, USA, 1968.
34. Chomsky, N. A Review of Skinner's Verbal Behavior. In *Readings in the Psychology of Language*; Jakobovits, L.A., Miron, M.S., Eds.; Prentice-Hall: Upper Saddle River, NJ, USA, 1967; pp. 142–143.
35. Kahneman, D. *Thinking, Fast and Slow*; Farrar, Straus and Giroux: New York, NY, USA, 2011.
36. Gigerenzer, G.; Goldstein, D.G. Reasoning the fast and frugal way: Models of bounded rationality. *Psychol. Rev.* **1996**, *103*, 650–669. [[CrossRef](#)] [[PubMed](#)]
37. Gigerenzer, G.; Todd, P.; Group, A. *Simple Heuristics that Make Us Smart*; Evolution and Cognition; Oxford University Press: Oxford, UK, 2000.
38. Gigerenzer, G.; Selten, R. *Bounded Rationality: The Adaptive Toolbox*; Dahlem Konferenzen; MIT Press: Cambridge, MA, USA, 2002.
39. Minsky, M.; Lee, J. *Society of Mind*; Touchstone Book; Simon & Schuster: New York, NY, USA, 1988.
40. Sanfey, A.G. Social Decision-Making: Insights from Game Theory and Neuroscience. *Science* **2007**, *318*, 598–602. [[CrossRef](#)] [[PubMed](#)]
41. Nash, J. Non-Cooperative Games (PhD Thesis). In *The Essential John Nash*; Kuhn, H., Nasar, S., Eds.; Princeton University Press: Princeton, NJ, USA, 2001; p. 244.
42. Von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behaviour*, 3rd ed.; Princeton University Press: Princeton, NJ, USA, 1974.
43. Friedman, M. *Essays in Positive Economics*; Phoenix Bks; University of Chicago Press: Chicago, IL, USA, 1953.

44. Shapley, D. Game Theorist Morgenstern Dies. *Science* **1977**, *197*, 649. [[CrossRef](#)] [[PubMed](#)]
45. Dreber, A.; Rand, D.G.; Fudenberg, D.; Nowak, M.A. Winners don't punish. *Nature* **2008**, *452*, 348–351. [[CrossRef](#)] [[PubMed](#)]
46. Nowak, M.A. Five Rules for the Evolution of Cooperation. *Science* **2006**, *314*, 1560–1563. [[CrossRef](#)] [[PubMed](#)]
47. Von Hippel, W.; Trivers, R. The evolution and psychology of self-deception. *Behav. Brain Sci.* **2011**, *34*, 1–16. [[CrossRef](#)] [[PubMed](#)]
48. Fehr, E.; Gächter, S. Altruistic punishment in humans. *Nature* **2002**, *415*, 137–140. [[CrossRef](#)] [[PubMed](#)]
49. Nowak, M.A.; Page, K.M.; Sigmund, K. Fairness Versus Reason in the Ultimatum Game. *Science* **2000**, *289*, 1773–1775. [[CrossRef](#)] [[PubMed](#)]
50. Axelrod, R.; Hamilton, W. The evolution of cooperation. *Science* **1981**, *211*, 1390–1396. [[CrossRef](#)] [[PubMed](#)]
51. Nowak, M.A.; Sigmund, K. Evolutionary Dynamics of Biological Games. *Science* **2004**, *303*, 793–799. [[CrossRef](#)] [[PubMed](#)]
52. Ball, P. When does punishment work? *Nature* **2009**. [[CrossRef](#)]
53. Gass, S.; Assad, A. *An Annotated Timeline of Operations Research: An Informal History*; International Series in Operations Research & Management Science; Kluwer Academic Publishers: Alphen aan den Rijn, The Netherlands, 2005.
54. Poundstone, W. *Prisoner's Dilemma*; Anchor Books; Knopf Doubleday: New York, NY, USA, 1992.
55. Axelrod, R. *The Evolution of Cooperation*; Basic Books: New York, NY, USA, 2006.
56. Nowak, M.A.; Sigmund, K. Game-dynamical aspects of the prisoner's dilemma. *Appl. Math. Comput.* **1989**, *30*, 191–213. [[CrossRef](#)]
57. Nowak, M.A.; May, R.M. Evolutionary games and spatial chaos. *Nature* **1992**, *359*, 826–829. [[CrossRef](#)]
58. Nowak, M.A.; Sigmund, K. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **1993**, *364*, 56–58. [[CrossRef](#)] [[PubMed](#)]
59. Ai, G.; Deng, Y. An overview of prospect theory in travel choice behavior under risk. In Proceedings of the 2017 4th International Conference on Transportation Information and Safety (ICTIS), Edmonton, AB, Canada, 8–10 August 2017; pp. 1128–1133.
60. Dörner, D. *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*; Merloyd Lawrence Book; Basic Books: New York, NY, USA, 1996.
61. Kahneman, D.; Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica* **1979**, *47*, 263–292. [[CrossRef](#)]
62. Kahnemann, D. Voice—10 Questions with... Daniel Kahneman on Humans and Decision Making. *J. Financ. Plan.* **2004**, *17*, 1013.
63. Gratch, J.; Marsella, S. A domain-independent framework for modeling emotion. *Cogn. Syst. Res.* **2004**, *5*, 269–306. [[CrossRef](#)]
64. Marsella, S.C.; Gratch, J.; Petta, P. Computational Models of Emotion. In *A Blueprint for an Affectively Competent Agent: Cross-Fertilization between Emotion Psychology, Affective Neuroscience, and Affective Computing*; Marsella, S., Gratch, J., Petta, P., Eds.; Oxford University Press: Oxford, UK, 2010.
65. Smith, A.M.; Lewis, C.; Hullet, K.; Smith, G.; Sullivan, A. An Inclusive View of Player Modeling. In Proceedings of the 6th FDG '11—International Conference on Foundations of Digital Games, Bordeaux, France, 28 June–1 July 2011; ACM: New York, NY, USA, 2011; pp. 301–303.
66. Yannakakis, G.N.; Spronck, P.; Loiacono, D.; André, E. Player Modeling. In *Artificial and Computational Intelligence in Games*; Lucas, S.M., Mateas, M., Preuss, M., Spronck, P., Togelius, J., Eds.; Dagstuhl Follow-Ups; Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik: Dagstuhl, Germany, 2013; Volume 6, pp. 45–59.
67. Gow, J.; Baumgarten, R.; Cairns, P.; Colton, S.; Miller, P. Unsupervised Modeling of Player Style with LDA. *IEEE Trans. Comput. Intell. AI Games* **2012**, *4*, 152–166. [[CrossRef](#)]
68. Simon, H.A. Motivational and Emotional Controls of Cognition. *Psychol. Rev.* **1967**, *74*, 29–39. [[CrossRef](#)] [[PubMed](#)]
69. Tversky, A.; Kahneman, D. The framing of decisions and the psychology of choice. *Science* **1981**, *211*, 453–458. [[CrossRef](#)] [[PubMed](#)]



70. De Martino, B.; Kumaran, D.; Seymour, B.; Dolan, R.J. Frames, Biases, and Rational Decision-Making in the Human Brain. *Science* **2006**, *313*, 684–687. [[CrossRef](#)] [[PubMed](#)]
71. Ajzen, I. From intentions to actions: A theory of planned behavior. In *Action Control*; Springer: Berlin/Heidelberg, Germany, 1985.
72. Ajzen, I. The theory of planned behavior. *Organ. Behav. Hum. Decis. Process.* **1991**, *50*, 179–211. [[CrossRef](#)]
73. Ajzen, I. Constructing a ToPB Questionnaire: Conceptual and Methodological Considerations. 2002. Available online: <http://www.unix.oit.umass.edu/tpb.measurement.pdf> (accessed on 13 March 2018).
74. Adamson, R. *Short History of Logic*; Forgotten Books: London, UK, 2009.
75. Restall, G. *Logic: An Introduction*; Routledge: Abingdon, UK, 1969.
76. Towell, A. The Psychology of Rational Thought: A Critical Estimate of Current Views and an Hypothesis Concerning the Role of Language in the Structure of Human Reason. Master's Thesis, University of British Columbia, Vancouver, BC, Canada, 1931.
77. Hastie, R.; Dawes, R. *Rational Choice in an Uncertain World: The Psychology of Judgment and Decision Making*; Sage: Newcastle upon Tyne, UK, 2009.
78. Hamilton, W.; Mansel, H.; Vietch, J. *Lectures on Metaphysics and Logic*; Number v. 1 in Lectures on Metaphysics and Logic; Gould & Lincoln: Boston, MA, USA, 1859.
79. Cohen, M.; Nagel, E.; Corcoran, J. *An Introduction to Logic*; Hackett Publishing Company: Indianapolis, IN, USA, 1993.
80. Kant, I. *Logic*; W. Simpkin and R. Marshall: London, UK, 1819.
81. Anderson, J. *Cognitive Psychology and Its Implications*; A Series of Books in Psychology; W.H. Freeman: New York, NY, USA, 1990.
82. Van Benthem, J.; van Ditmarsch, H.; Ketting, J.; Meyer-Viol, W. *Logica voor Informatici*; Addison-Wesley Nederland: Amsterdam, The Netherlands, 1991.
83. Parry, W.T.; Hacker, E.A. *Aristotelian Logic*; State University of New York Press: Albany, NY, USA, 1991.
84. Priest, G. *Logic: A Very Short Introduction*; Very Short Introductions; Oxford University Press: Oxford, UK, 2000.
85. Blackburn, P.; deRijke, M.; Venema, Y. *Modal Logic*; Cambridge U. Press: Cambridge, UK, 2001.
86. Hamilton, A. *Logic for Mathematicians*; Cambridge University Press: Cambridge, UK, 1988.
87. Joyner, D. *Adventures in Group Theory: Rubik's Cube, Merlin's Machine, and Other Mathematical Toys*; Adventures in Group Theory; Johns Hopkins University Press: Baltimore, MD, USA, 2008.
88. Goldblatt, R. Mathematical modal logic: A view of its evolution. *J. Appl. Logic* **2003**, *1*, 309–392. [[CrossRef](#)]
89. Kripke, S.A. A Completeness Theorem in Modal Logic. *J. Symb. Logic* **1959**, *24*, 1–14. [[CrossRef](#)]
90. Hansson, S.O. Semantics for more plausible deontic logics. *J. Appl. Logic* **2004**, *2*, 3–18. [[CrossRef](#)]
91. Post, E.L. Introduction to a General Theory of Elementary Propositions. *Am. J. Math.* **1921**, *43*, 163–185. [[CrossRef](#)]
92. Gilbert, D. Buried by bad decisions. *Nature* **2011**, *474*, 275–277. [[CrossRef](#)] [[PubMed](#)]
93. Davis, M. *Game Theory: A Nontechnical Introduction*; Dover Books on Mathematics; Dover Publications: Mineola, NY, USA, 1997.
94. Osborne, M.J.; Rubenstein, A. *A Course in Game Theory*; The MIT Press: Cambridge, MA, USA, 1994.
95. Hilbe, C.; Traulsen, A. Emergence of responsible sanctions without second order free riders, antisocial punishment or spite. *Sci. Rep.* **2012**, *2*, 458. [[CrossRef](#)] [[PubMed](#)]
96. Tversky, A.; Kahneman, D. Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncert.* **1992**, *5*, 297–323. [[CrossRef](#)]
97. Gibson, D.R. Nuclear deterrence: Decisions at the brink. *Nature* **2012**, *487*, 27–29. [[CrossRef](#)] [[PubMed](#)]
98. Sagan, S.D. Policy: A call for global nuclear disarmament. *Nature* **2012**, *487*, 30–32. [[CrossRef](#)] [[PubMed](#)]
99. Schopenhauer, A. *Die Welt als Wille und Vorstellung*; Number 1 in Die Welt als Wille und Vorstellung, Brockhaus; Library of Alexandria: Alexandria, Egypt, 1844.
100. Kahneman, D. Maps of Bounded Rationality: Psychology for Behavioral Economics. *Am. Econ. Rev.* **2003**, *93*, 1449–1475. [[CrossRef](#)]
101. Edwards, W. The theory of decision making. *Psychol. Bull.* **1954**, *41*, 380–417. [[CrossRef](#)]
102. Edwards, W.; Shanteau, J.; Mellers, B.; Schum, D. *Decision Science and Technology: Reflections on the Contributions of Ward Edwards*; Kluwer Academic: Dordrecht, The Netherlands, 1999.

103. Luce, R.; Raiffa, H. *Games and Decisions: Introduction and Critical Survey*; Dover Books on Mathematics; Dover Publications: Mineola, NY, USA, 1989.
104. Luce, R. *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*; Scientific Psychology Series; Lawrence Erlbaum: Mahwah, NJ, USA, 2000.
105. Barberis, N.C. Thirty Years of Prospect Theory in Economics: A Review and Assessment. *J. Econ. Perspect.* **2013**, *27*, 173–196. [[CrossRef](#)]
106. Hey, J.D.; Lotito, G.; Maffioletti, A. *The Descriptive and Predictive Adequacy of Theories of Decision Making under Uncertainty/Ambiguity*; Discussion Papers 08/04; Department of Economics, University of York: York, UK, 2008.
107. Luce, R.D. Where Does Subjective Expected Utility Fail Descriptively? *J. Risk Uncert.* **1992**, *5*, 5–27.
108. Brown, S. *Play as an Organizing Principle: Clinical Evidence and Personal Observations*; Cambridge University Press: Cambridge, UK, 1998; Chapter 12, pp. 243–260.
109. Bruce, T. *Developing Learning in Early Childhood*; 0–8 Series; Paul Chapman: London, UK, 2004.
110. Pasin, F.; Giroux, H. The impact of a simulation game on operations management education. *Comput. Educ.* **2011**, *57*, 1240–1254. [[CrossRef](#)]
111. Christoph, N. The Role of Metacognitive Skills in Learning to Solve Problems. Ph.D. Thesis, Universiteit van Amsterdam, Amsterdam, The Netherlands, 2006.
112. Pee, N.C. Computer Games Use in an Educational System. Ph.D. Thesis, University of Nottingham, Nottingham, UK, 2011.
113. Sandberg, J.; Christoph, N.; Emans, B. Tutor training: A systematic investigation of tutor requirements and an evaluation of a training. *Br. J. Educ. Technol.* **2001**, *32*, 69–90. [[CrossRef](#)]
114. Warren, S.J.; Jones, G.; Trombley, A. Skipping Pong and moving straight to World of Warcraft: The challenge of research with complex games used for learning. *Int. J. Web Based Communities* **2011**, *7*, 218–233. [[CrossRef](#)]
115. Malone, T.; Lepper, M.R. Making learning fun: A Taxonomy of Intrinsic Motivations for Learning. *Aptit. Learn. Instruct.* **1987**, *3*, 223–235.
116. Schneider, M.; Carley, K.M.; Moon, I.C. *Detailed Comparison of America's Army Game and Unit of Action Experiments*; Technical Report; Carnegie Mellon University: Pittsburgh, PA, USA, 2005.
117. Habgood, M.P.J. The Effective Integration of Digital Games and Learning Content. Ph.D. Thesis, University of Nottingham, Nottingham, UK, 2007.
118. Squire, K.; Barnett, M.; Grant, J.M.; Higginbotham, T. Electromagnetism supercharged!: Learning physics with digital simulation games. In Proceedings of the 6th International Conference on Learning Sciences, Santa Monica, CA, USA, 22–26 June 2004; ACM Press: New York, NY, USA, 2004; pp. 513–520.
119. Young, J.; Uptis, R. The microworld of Phoenix Quest: Social and cognitive considerations. *Educ. Inf. Technol.* **1999**, *4*, 391–408. [[CrossRef](#)]
120. Ford, C.W., Jr.; Minsker, S. TREEZ—An educational data structures game. *J. Comput. Sci. Coll.* **2003**, *18*, 180–185.
121. Zhu, Q.; Wang, T.; Tan, S. Adapting Game Technology to Support Software Engineering Process Teaching: From SimSE to MO-SEProcess. In Proceedings of the ICNC '07 Third International Conference on Natural Computation, Haikou, China, 24–27 August 2007; IEEE Computer Society: Washington, DC, USA, 2007; Volume 5, pp. 777–780.
122. Beale, I.L.; Kato, P.M.; Marin-Bowling, V.M.; Guthrie, N.; Cole, S.W. Improvement in Cancer-Related Knowledge Following Use of a Psychoeducational Video Game for Adolescents and Young Adults with Cancer. *J. Adolesc. Health* **2007**, *41*, 263–270. [[CrossRef](#)] [[PubMed](#)]
123. Lennon, J.L. Debriefings of web-based malaria games. *Simul. Gaming* **2006**, *37*, 350–356. [[CrossRef](#)]
124. Roubidoux, M.A. Breast cancer detective: A computer game to teach breast cancer screening to Native American patients. *J. Cancer Educ.* **2005**, *20*, 87–91. [[CrossRef](#)] [[PubMed](#)]
125. Connolly, T.; Boyle, E.; Stansfield, M.; Hainey, T. A Survey of Students' computer game playing habits. *J. Adv. Technol. Learn.* **2007**, *4*, 218–223. [[CrossRef](#)]
126. Connolly, T.; Boyle, E.; Hainey, T. A Survey of Students' Motivations for Playing Computer Games: A Comparative Analysis. In Proceedings of the 1st European Conference on Games-Based Learning (ECGBL), Scotland, UK, 25–26 October 2007.

127. Puschel, T.; Lang, F.; Bodenstein, C.; Neumann, D. A Service Request Acceptance Model for Revenue Optimization—Evaluating Policies Using a Web Based Resource Management Game. In Proceedings of the HICSS '10 2010 43rd Hawaii International Conference on System Sciences, Kauai, HI, USA, 1–8 January 2010; IEEE Computer Society: Washington, DC, USA, 2010; pp. 1–10.
128. Babb, E.M.; Leslie, M.A.; Slyke, M.D.V. The Potential of Business-Gaming Methods in Research. *J. Bus.* **1966**, *39*, 465–472. [[CrossRef](#)]
129. Rowland, K.M.; Gardner, D.M. *The Uses of Business Gaming in Education and Laboratory Research*; Vol. BEBR No. 10; Education, Gaming; University of Illinois: Champaign, IL, USA, 1971.
130. Connolly, T.; Stansfield, M.; Josephson, J.; L'azaro, N.; Rubio, G.; Ortiz, C.; Tsvetkova, N.; Tsvetanova, S. Using Alternate Reality Games to Support Language Learning. In Proceedings of the Web-Based Education WBE 2008, Innsbruck, Austria, 17–19 March 2008.
131. Peschon, J.; Isaksen, L.; Tyler, B. The growth, accretion, and decay of cities. In Proceedings of the 1996 International Symposium on Technology and Society Technical Expertise and Public Decisions, Princeton, NJ, USA, 21–22 June 1996; pp. 301–310.
132. Kuemmerle, W. Launching a High-Risk Business—An Interactive Simulation. Developed by High Performance Systems, Inc., based on the work of William A. Sahlman and Michael J. Roberts. *Small Bus. Econ.* **2000**, *15*, 243–245. [[CrossRef](#)]
133. Gee, J.P. *What Video Games Have to Teach Us About Learning and Literacy*; Palgrave Macmillan: Basingstoke, UK, 2003.
134. Genesereth, M.; Love, N. General game playing: Overview of the AAAI competition. *AI Mag.* **2005**, *26*, 62–72.
135. Love, N.; Hinrichs, T.; Haley, D.; Schkufza, E.; Genesereth, M. *General Game Playing: Game Description Language Specification*; Technical Report; Stanford Logic Group: Stanford, CA, USA, 2008.
136. Blackburn, P.; van Benthem, J.; Wolter, F. *Handbook of Modal Logic*; Elsevier: New York, NY, USA, 2005.
137. Fagin, R.; Halpern, J.Y.; Moses, Y.; Vardi, M.Y. *Reasoning about Knowledge*; MIT Press: Cambridge, MA, USA, 1995.
138. Hildmann, H.; Hildmann, J. A formalism to define, assess and evaluate player behaviour in mobile device based serious games. In *Serious Games and Edutainment Applications*; Springer: London, UK, 2012; Chapter 6.
139. Shankar, N. *Cambridge Tracts in Theoretical Computer Science 38: Metamathematics, Machines, and Gödel's Proof*; Cambridge University Press: Cambridge, UK, 1994.
140. Pizzi, D.; Cavazza, M.; Whittaker, A.; Lugin, J.L. Automatic Generation of Game Level Solutions as Storyboards. In Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE), Stanford, CA, USA, 22–24 October 2008; The AAAI Press: Palo Alto, CA, USA, 2008.
141. Ochert, A. Nature versus nurture at 60 MHz. *Nature* **1997**, *385*, 499–499. [[CrossRef](#)]
142. Fishbein, M.; Ajzen, I.; Albarracín, D.; Hornik, R. *Prediction and Change of Health Behavior: Applying the Reasoned Action Approach*; L. Erlbaum Associates: Mahwah, NJ, USA, 2007.
143. Cowley, B.; Charles, D.; Black, M.; Hickey, R. Analyzing player behavior in pacman using feature-driven decision theoretic predictive modeling. In Proceedings of the CIG'09 5th international conference on Computational Intelligence and Games, Milano, Italy, 7–10 September 2009; IEEE Press: Piscataway, NJ, USA, 2009; pp. 170–177.
144. Hildmann, H. Behavioural game AI—A theoretical approach. In Proceedings of the 2011 International Conference for Internet Technology and Secured Transactions (ICITST), Abu Dhabi, UAE, 11–14 December 2011; pp. 550–555.
145. Bitterberg, T.; Hildmann, H.; Branki, C. Using resource management games for mobile phones to teach social behaviour. In *Proceeding of 2008 Conference on Techniques and Applications for Mobile Commerce*; IOS Press: Amsterdam, The Netherlands, 2008; pp. 77–84.
146. Hildmann, J.; Hildmann, H. *Computer Games in Education*; Encyclopedia of Computer Graphics and Games; Springer International Publishing: Cham, Switzerland, 2018; Chapter 278.
147. Hildmann, H.; Livingstone, D. A Formal Approach to Represent, Implement and Assess Learning Targets in Computer Games. In Proceedings of the 1st European Conference on Games-Based Learning (ECGBL), Scotland, UK, 25–26 October 2007; Academic Conferences, Limited: Paisley, UK, 2007.

148. Hildmann, H.; Hainey, T.; Livingstone, D. Psychology and logic: Design considerations for a customisable educational resource management game. In Proceedings of the 5th Annual International Conference in Computer Game Design and Technology, Liverpool, UK, 2007. Available online: [https://s3.amazonaws.com/academia.edu.documents/30800375/SP1.pdf?AWSAccessKeyId=AKIAIWOWYYGZ2Y53UL3A&Expires=1537866177&Signature=4JAas10rDHnfMz98vA30nDFNyPI%3D&response-content-disposition=inline%3B%20filename%3DPsychology\\_and\\_logic\\_design\\_consideratio.pdf](https://s3.amazonaws.com/academia.edu.documents/30800375/SP1.pdf?AWSAccessKeyId=AKIAIWOWYYGZ2Y53UL3A&Expires=1537866177&Signature=4JAas10rDHnfMz98vA30nDFNyPI%3D&response-content-disposition=inline%3B%20filename%3DPsychology_and_logic_design_consideratio.pdf) (accessed on 25 September 2018)
149. Hildmann, H.; Branki, C.; Pardavila, C.; Livingstone, D. A framework for the development, design and deployment of customisable mobile and hand held device based serious games. In Proceedings of the 2nd European Conference on Games Based Learning (ECGBL08), Barcelona, Spain, 16 October 2008; API, Academic Publishing International: Barcelona, Spain, 2008.



© 2018 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).